



US 20180183901A1

(19) **United States**

(12) **Patent Application Publication**
Lariviere et al.

(10) **Pub. No.: US 2018/0183901 A1**

(43) **Pub. Date: Jun. 28, 2018**

(54) **MESSAGE PROCESSING PROTOCOL
WHICH MITIGATES OPTIMISTIC
MESSAGING BEHAVIOR**

G06Q 40/02 (2012.01)

G06Q 40/06 (2012.01)

G06Q 40/04 (2012.01)

(52) **U.S. Cl.**

CPC **H04L 69/08** (2013.01); **G06Q 40/08**
(2013.01); **G06Q 40/04** (2013.01); **G06Q**
40/06 (2013.01); **G06Q 40/02** (2013.01)

(71) Applicant: **Chicago Mercantile Exchange Inc.,**
Chicago, IL (US)

(72) Inventors: **David Alan Lariviere**, Chicago, IL
(US); **Bernard Pieter Hosman**,
Amsterdam (NL); **Pearce Ian**
Peck-Walden, Chicago, IL (US); **Ari L.**
Studnitzer, Northbrook, IL (US);
Zachary Bonig, Skokie, IL (US);
Manmathasivaram Nagarajan,
Chicago, IL (US)

(57)

ABSTRACT

The disclosed embodiments relate to implementation, such as by a message processing system architecture, of a message handling system and/or protocol which mitigates optimistic messaging behavior. As used herein optimistic messaging behavior may at least refer to the transmission, in whole or in part, of a message, or of one or more messages of a sequence thereof, to a receiving system, such as the system implemented by the disclosed embodiments, prior to the sender determining that the message(s) is/are desired, alone or in conjunction with subsequently canceling or otherwise invalidating the message(s) once it has been determined that it is not desired. The disclosed embodiments mitigate such behavior by detecting and taking action with respect to these types of messages to deter optimizing behavior which may or may not be manipulative, while maintaining reliable message handling under increasing processing loads with minimal impact on users being able to send messages to the system.

(21) Appl. No.: **15/851,067**

(22) Filed: **Dec. 21, 2017**

Related U.S. Application Data

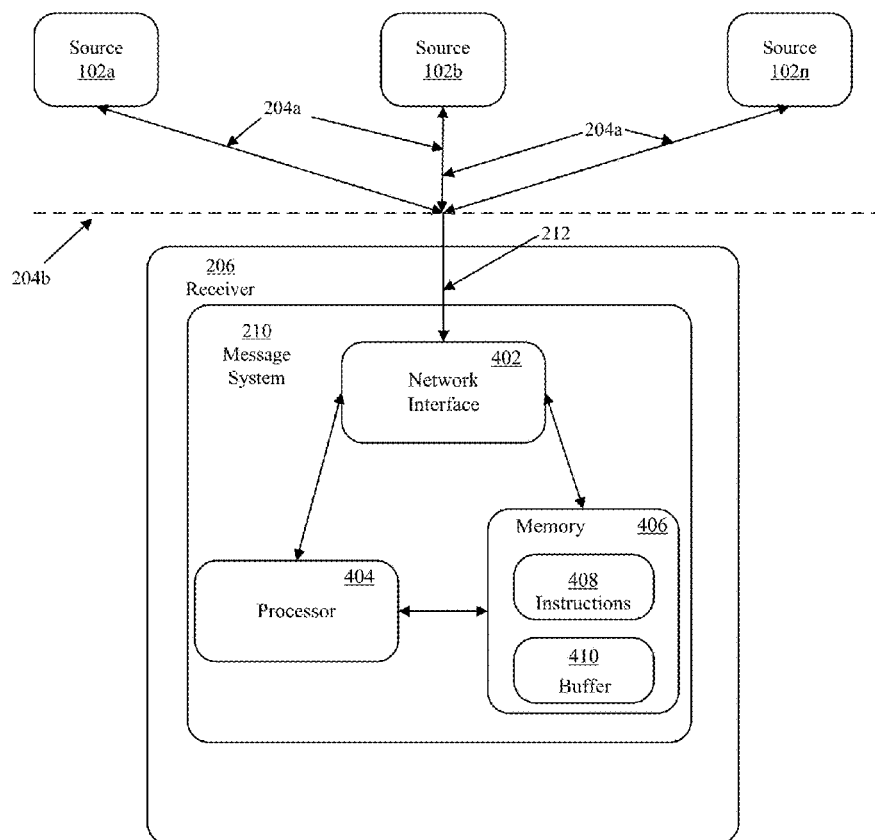
(60) Provisional application No. 62/439,202, filed on Dec. 27, 2016.

Publication Classification

(51) **Int. Cl.**

H04L 29/06 (2006.01)

G06Q 40/08 (2012.01)



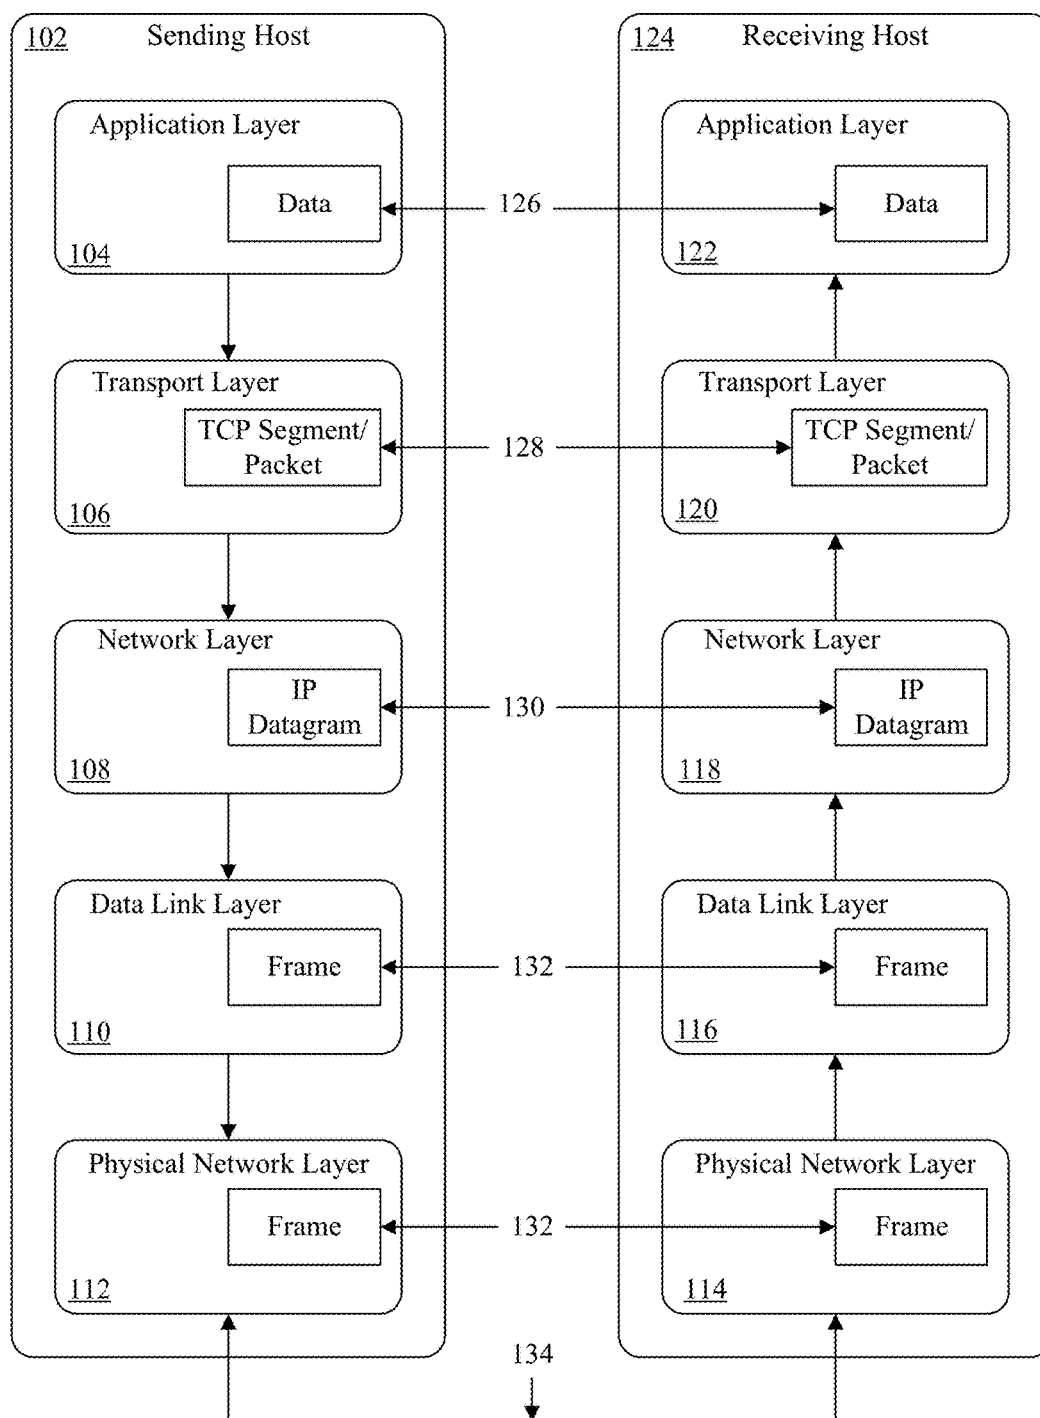


FIG. 1

FIG. 2




FIG. 3A





FIG. 3B



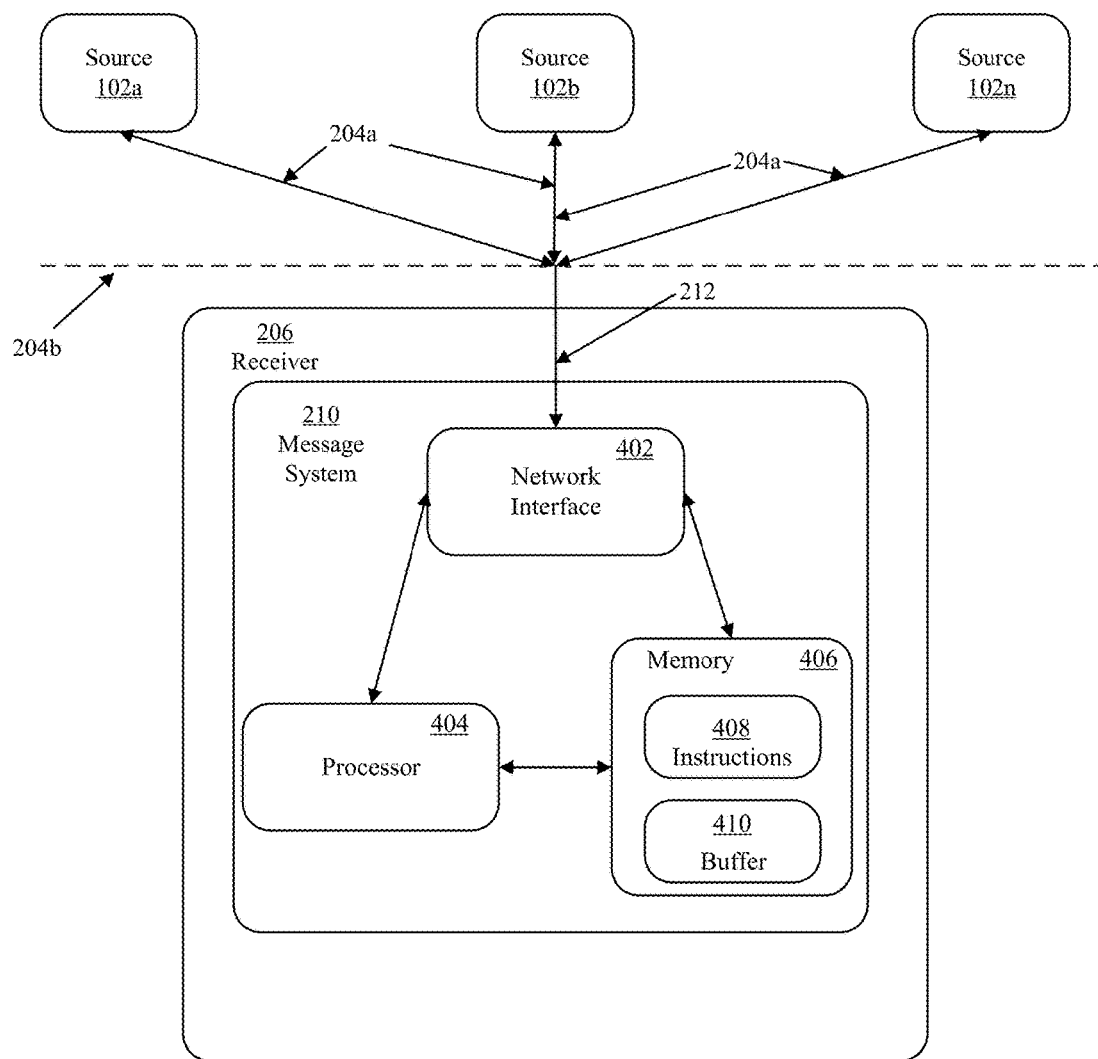


FIG. 4A

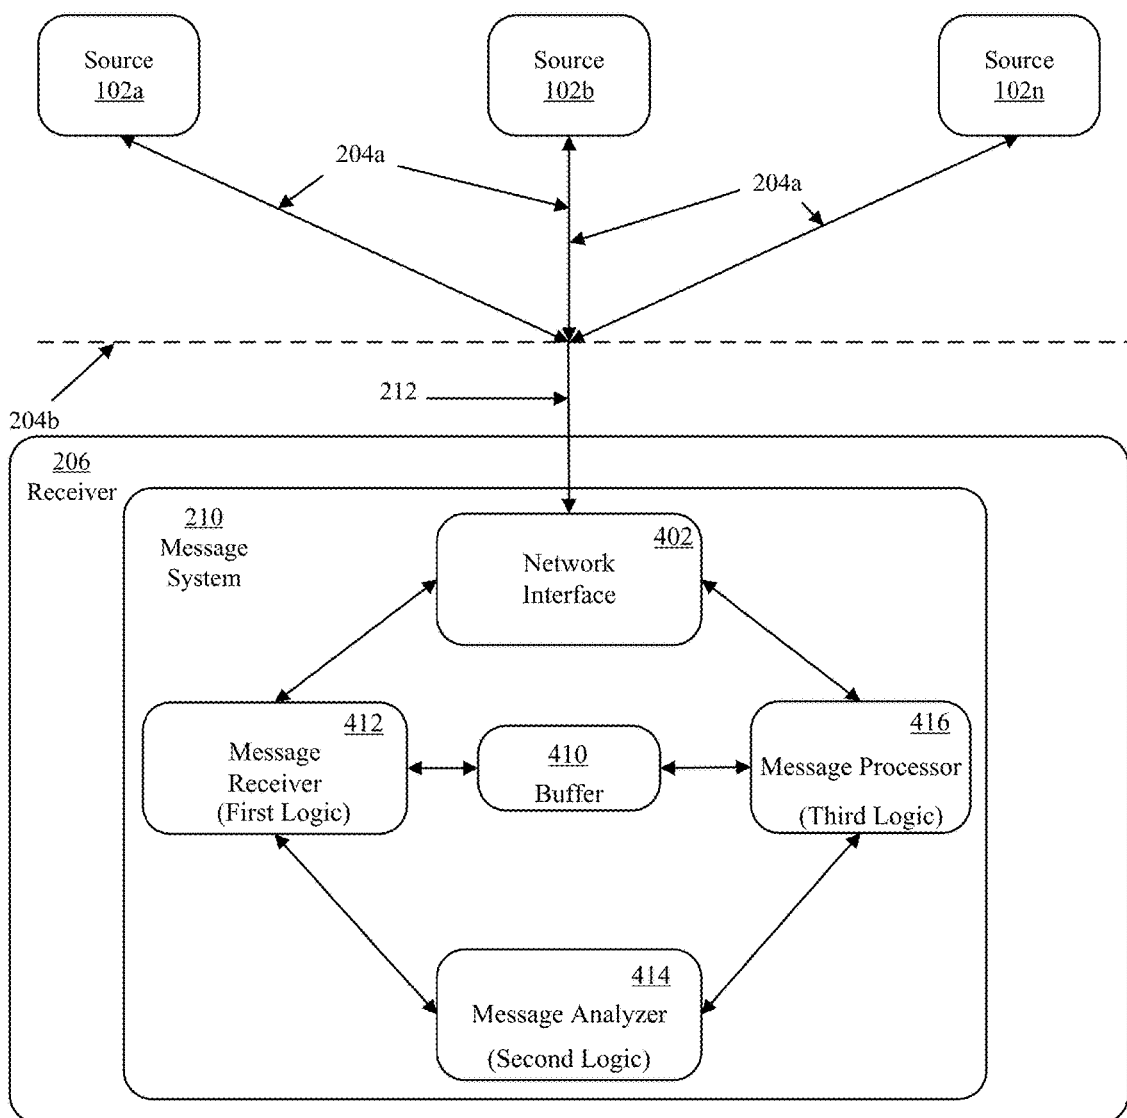


FIG. 4B

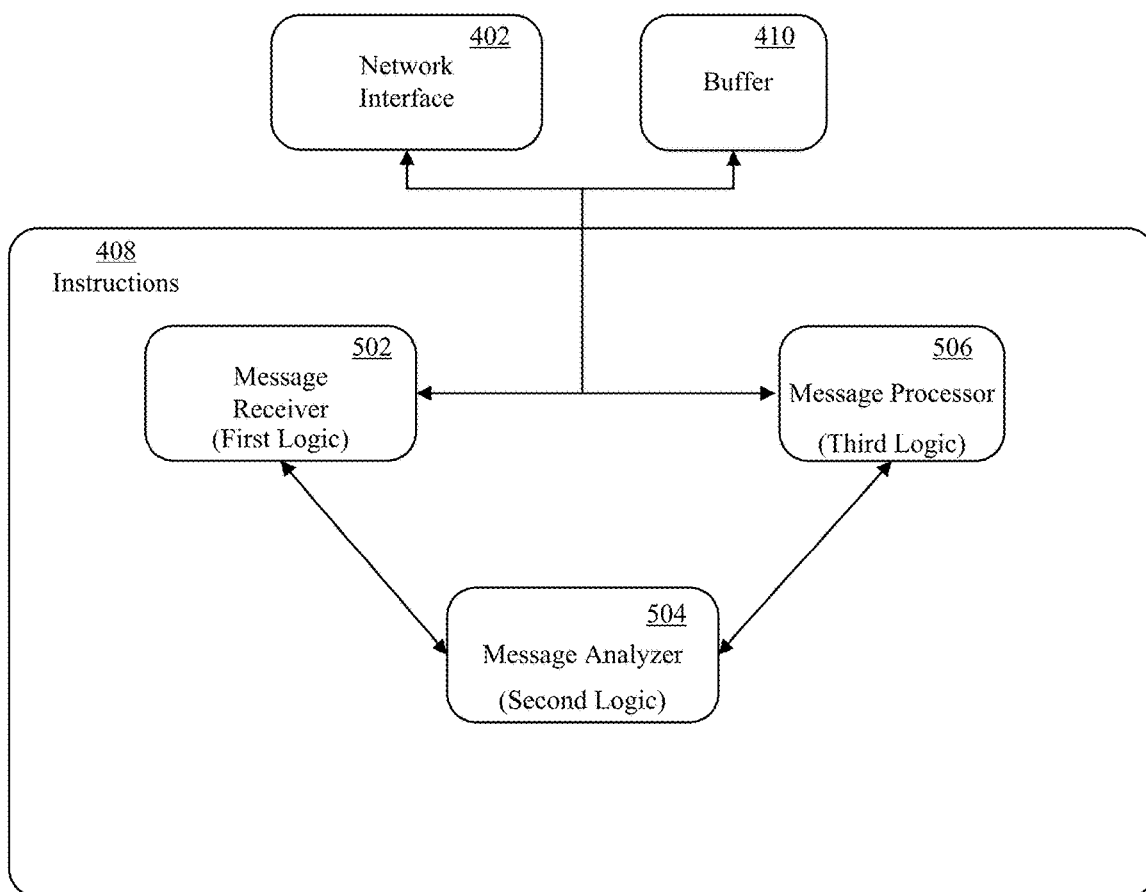


FIG. 5

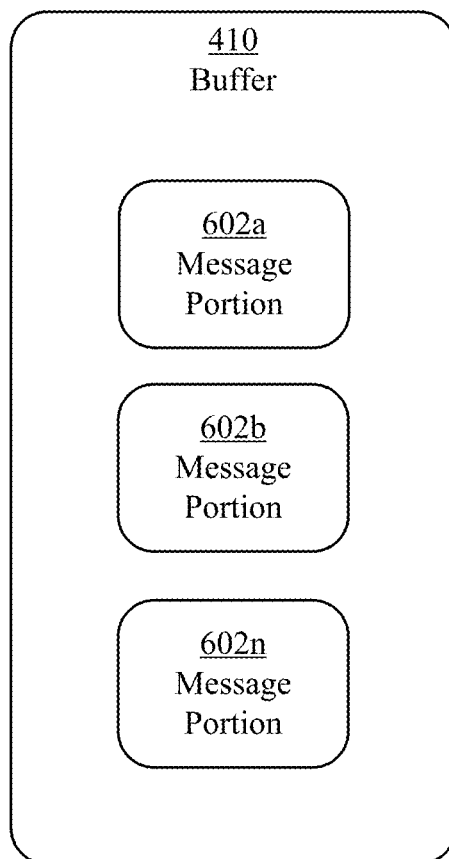


FIG. 6

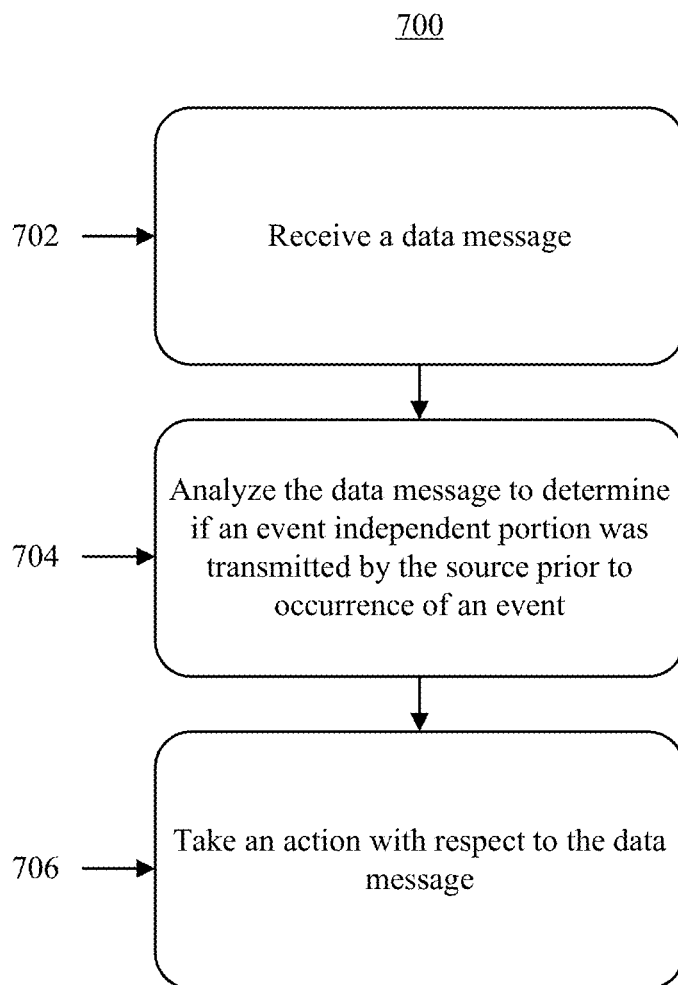


FIG. 7

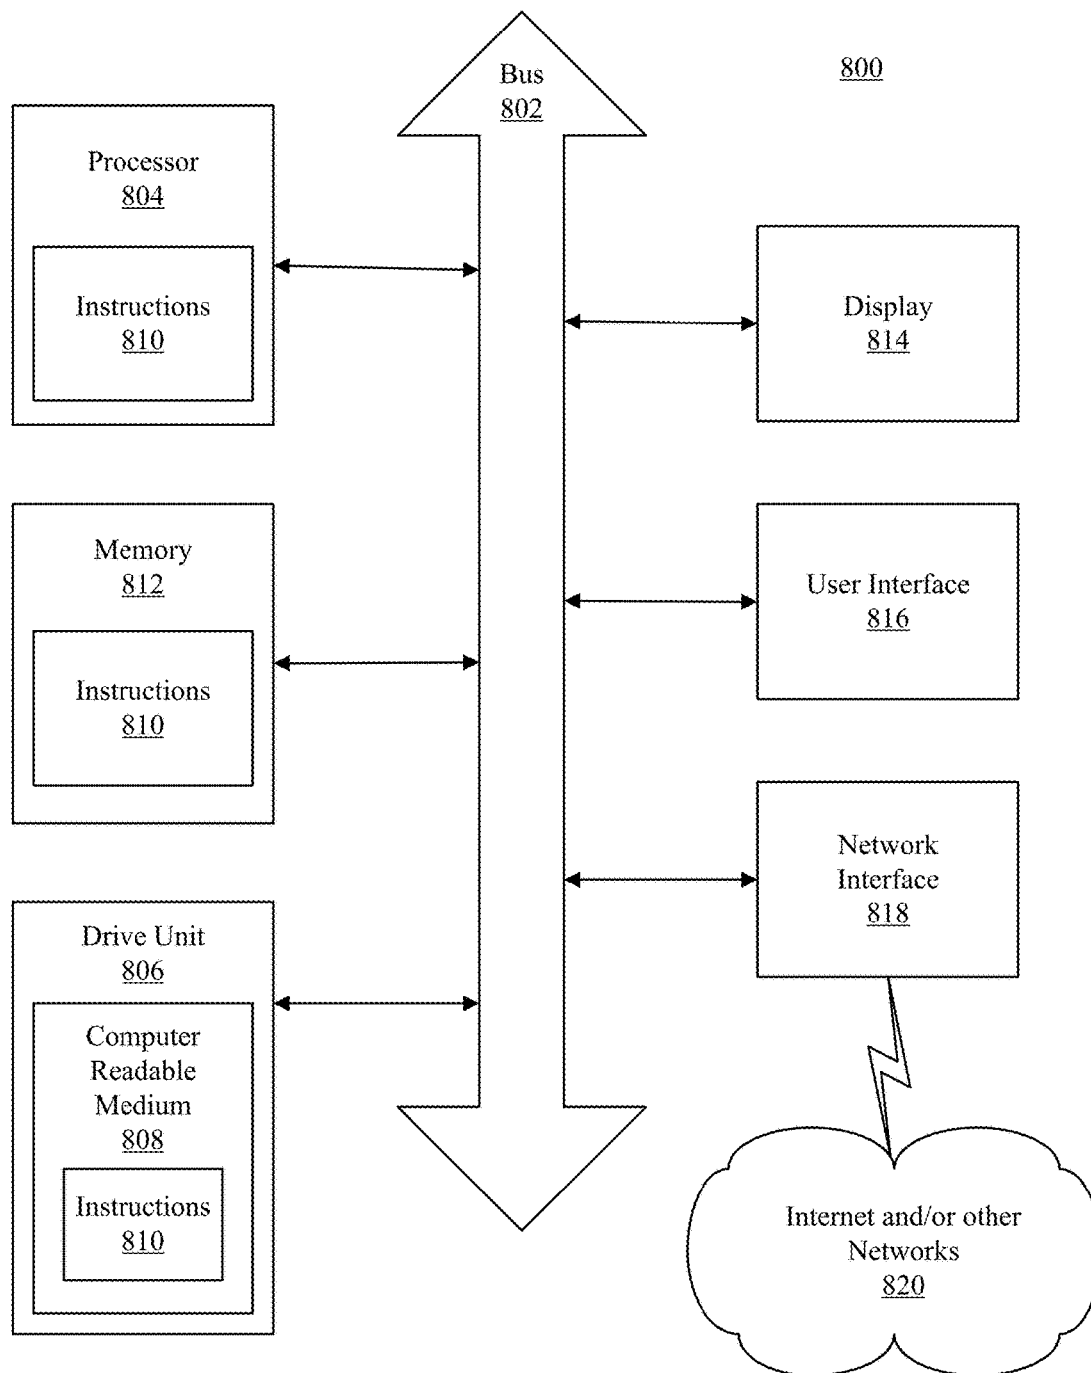


FIG. 8

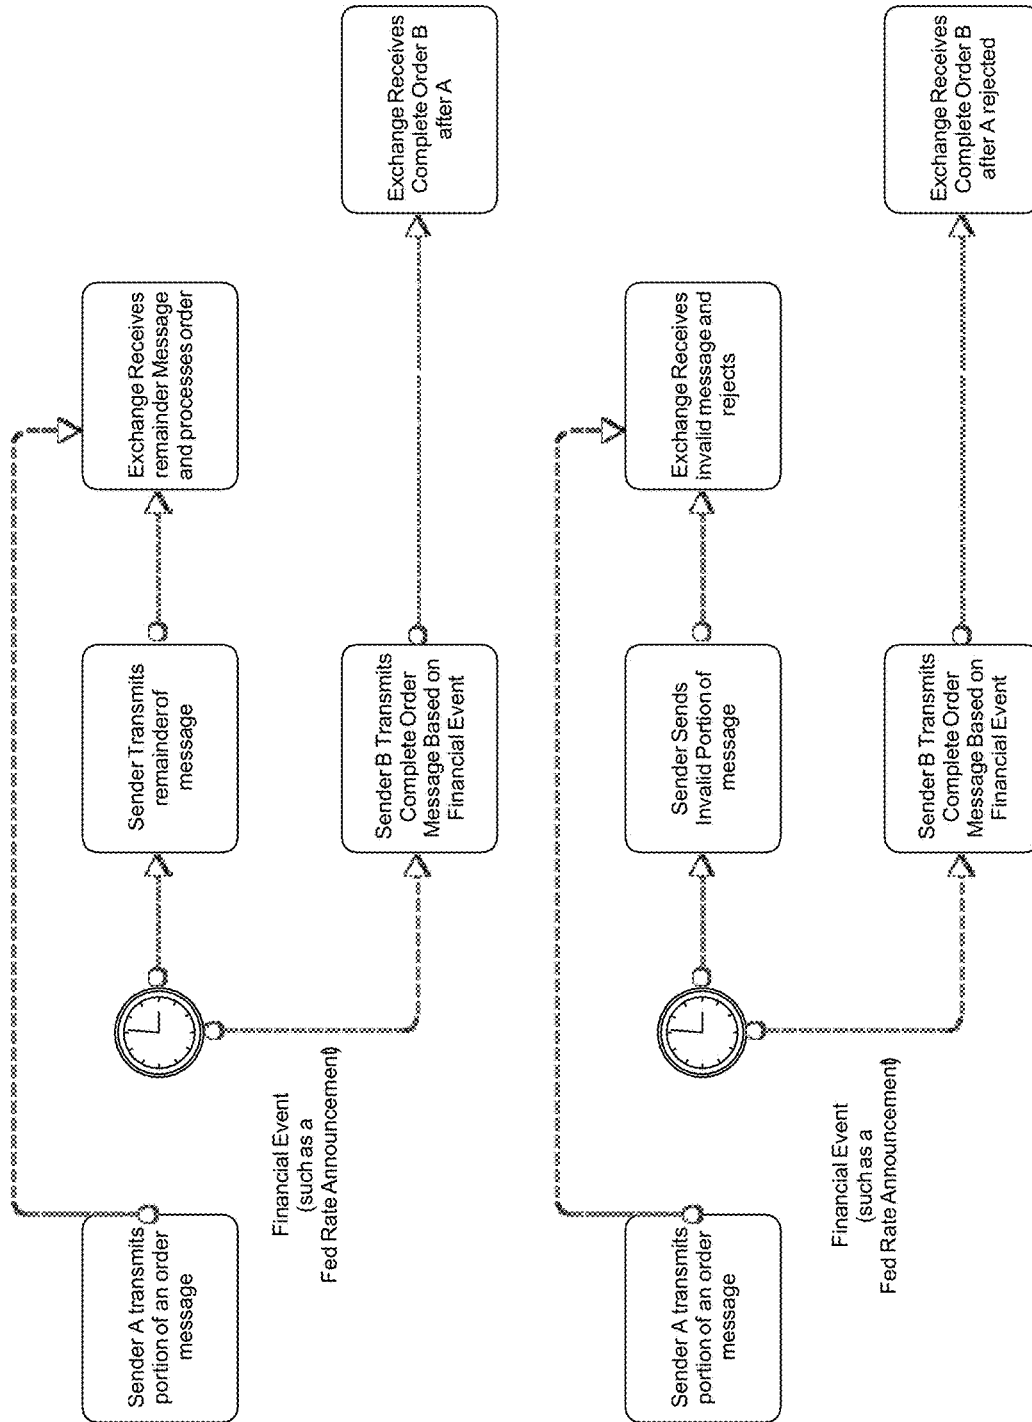


FIG. 9

**MESSAGE PROCESSING PROTOCOL
WHICH MITIGATES OPTIMISTIC
MESSAGING BEHAVIOR**

REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of the filing date under 35 U.S.C. § 119(e) of U.S. Provisional Application Ser. No. 62/439,202 filed Dec. 27, 2016 which is hereby incorporated by reference in its entirety.

BACKGROUND

[0002] Transaction processing systems may rely upon messages which are exchanged or otherwise communicated, such as via a network, e.g. the Internet, between the transaction processing system and the users of that system, and/or between other systems, to facilitate/implement the functionality of the transaction processing system. This exchange of messages may utilize particular communication protocols which may be standardized and well publicized so as to ease use of the systems, to ensure compatibility with other systems and communications media, and/or to facilitate use, development of and/or interaction with third party components. Multiple types of messaging protocols may be used, in combination and/or for different purposes. These protocols may be further defined to provide for efficient use of the communication medium, support for varying types and/or sizes of messages, secure transmission and/or reliable delivery of the messages communicated therewith.

[0003] A message protocol, also referred to as a communication protocol, generally specifies a set of rules and descriptions that detail how data should be structured, e.g. logically and/or physically, to allow for, or otherwise govern, transmission over a communications medium, across a network and/or between or within devices. Communication protocols exist that allow for the transmission of many different types of information and typically define or otherwise specify how the data should be structured, how to detect errors in the transmission of the data, the rate, or range thereof, of transmission for the data, how the data is transmitted, as well as other aspects of the mechanics of communication or combinations thereof.

[0004] To facilitate communication between computer systems, the communications protocols used thereby may be implemented according to the Open Systems Interconnection ("OSI") reference model. The OSI model conceptualizes a layered abstraction structure to handle the packaging and transmission of data between/across various network devices. The OSI model generally specifies seven different layers, each of which may also be referred to as an abstraction layer and the collection of layers may be referred to as a protocol stack, where each layer has a particular set of communication protocols which will work on that particular layer and allow for communication with the adjacent layers. It will be appreciated that the terms "layer" and "protocol stack" may also refer to the computer software programs/components which implement the various model layers. These layers include a physical layer, a data link layer, a network layer, a transport layer, a session layer, a presentation layer, and an application layer. It will be appreciated that the layered structure may be implementation dependent and that fewer or more layers may be defined, e.g. some layers may not be implemented, the functionality of multiple layers

may be combined into a single layer or the functionality of a single layer may be further divided among additional/multiple layers.

[0005] According to the OSI model, the physical layer deals with the electrical, or optical, and physical requirements for devices engaged in data transmission. The data link layer is the protocol layer that transfers data between adjacent network nodes in a wide area network (WAN) or between nodes on the same local area network (LAN) segment and provides the functional and procedural means to transfer data between network entities. The network layer provides the functional and procedural means of transferring fixed or variable length data sequences from one node to another connected to the same network. The transport layer provides the functional and procedural means of transferring variable-length data sequences from a source to a destination host via one or more networks. The session layer controls the dialogues between computers by establishing, managing, and terminating the connections, i.e. sessions, between local and remote applications. The presentation layer establishes context between application-layer entities, in which the application-layer entities may use different syntax and semantics if the presentation service provides a mapping between them. The application layer is the OSI layer closest to the end user and interacts with software applications on a user computer, which means both the OSI application layer and the user may interact directly with a software application. Among other attributes, the OSI model, and accompanying protocols compliant therewith, facilitates both unreliable and reliable communication, i.e. protocols which may or may not include mechanisms that detect and notify a user or other component if the delivery of data fails.

[0006] An example protocol that is used to implement the transport layer of the OSI model is the Transmission Control Protocol (TCP). TCP is known as a reliable protocol and is used in conjunction with the Internet Protocol (IP), e.g. IPv4 or IPv6, a typical protocol used to implement the network layer. According to TCP, data transmitted between computers is organized into segments, which are often referred to as packets, each of which identify the source of the data being transmitted and the destinations thereof, as well as may include a data payload comprising the data being transmitted, or a portion thereof. Each packet includes a packet header and, optionally, a payload. TCP is considered a reliable protocol because it uses sequence numbers to identify each packet of data that is sent from one computer to another computer so that all of the packets may be accounted for, i.e. the sequence numbers assigned to each packet identify the offset within a stream of data for which a particular packet contains data such that lost data may be caused to be retransmitted. The sequence number identifies the order of packets sent from each computer so that the data can be reconstructed in order, regardless of any packet reordering, or packet loss, that may occur during transmission from device to device as the packets traverse the network infrastructure. TCP uses acknowledgements sent by the receiver of data to tell the sender of that data that the data has been received by the receiver. In the event of packet loss, TCP allows for implementers of the protocol to choose how packet loss is handled. In one implementation, the receiver of the packets who detects packet loss may request that the sender resend only the missing packet(s). In another implementation, the receiver of the packets who detects packet loss may request that the sender resend all of the packets

including the missing packet(s) and packets already received from a particular demarcation point, e.g. sent since the opening of a connection or sent since a loss was detected. Additionally, TCP uses IP to effectuate the actual delivery of the data packets that are sent by each computer over a network. The combination of TCP/IP allows for transmission of data over a network, such as the internet, in a reliable manner.

[0007] An example of a network that uses TCP/IP to exchange data packets, or simply packets, between nodes on the network is a packet switched network. A packet switched network is a network that employs packet switching, a transmission process that involves breaking up a data message into multiple packets and transmitting each packet along the network from a source to a destination. The path that the packets take along the network may cover multiple different nodes on the network. The packets do not need to travel the same path from source to destination. For example, given a message consisting of packets 1, 2, 3, and 4 with a source node of A and a destination node of E for the message, packet 1 may travel from node A to node B to node E, whereas packet 2 may travel from node A to node C to node E, and so on. An example of a packet switched network is the Internet, or a Local Area Network (LAN.) Using TCP, as described herein, whether the packets traverse the same path from the sender to the receiver or not, and/or whether packets are delivered out of order due to inequities in the latency of different paths utilized, is irrelevant as the TCP protocol ensures delivery and facilitates reordering of the packets upon receipt into the order as transmitted.

[0008] The TCP protocol is a reliable protocol, meaning that the TCP protocol ensures the delivery of all packets. Part of how the TCP protocol ensures delivery of all packets is by using a three way handshake to establish a connection. The three way handshake involves a sending node sending a message to a receiving node that the sending node is ready to transmit data and wants to open a logical connection with the receiving node. The receiving node receives the request to open a connection and sends its own acknowledgement of the request. The sender and receiver then begin transmitting along the connection, each using an initial sequence number for their side of the connection. All information that is transmitted is broken up into segments, and each segment has a timer associated with it. If a sender does not receive acknowledgment that a segment has been received before the timer for that segment runs out, the sender resends the segment. On the receiving end, the receiver of the segments will reorder segments into the proper order if need be, because, for example, some segments timed out and had to be retransmitted out of order, or the segments were transmitted via different network paths having varying latency allowing some segments to overtake other segments to arrive first.

[0009] User Datagram Protocol (UDP) is as another Transport Layer protocol, which is an alternative to TCP. UDP is not, as opposed to TCP, a reliable protocol. UDP is connectionless, and makes a best effort for transmission. UDP does not use a handshake system between sender and receiver. UDP is typically used for systems where dropped packets are preferable to waiting for delayed packets, or segments of data.

[0010] A lower layer down from the Transport and Network Layers is the Data Link Layer which, as described above, defines the protocol to establish and terminate a

connection/session between two physically connected devices. Ethernet is an example of a data link layer protocol that can work in conjunction with the protocols which implement the other OSI layers, such as the TCP/IP protocols used for the Transport and Network layers. Ethernet describes how networked devices can format data for transmission to other network devices on the same network. Ethernet organizes the data to be transmitted into packets, referred to as Ethernet Frames, each of which identify the Ethernet source of data and Ethernet destination for the data, as well as include a data payload carrying, depending upon the amount, all or a subset of the data to be transmitted. An Ethernet Frame comprises a source and destination address, where the source and destination addresses correspond to Media Access Control addresses (or MAC addresses) for nodes on the network. The data payload is where any other headers for other protocols may be located as well as any data that may ultimately be used by applications. A TCP packet, as described above, may be transmitted over an Ethernet compatible network using one or more Ethernet frames.

[0011] The Application Layer of the OSI Model is the layer of the OSI model closest to an end user, and is where applications and end user processes are implemented that deal with the data transmitted from the bottom layers to the top layers. Some example protocols that handle data on the Application Layer are HTTP, SMTP, etc. World Wide Web browser programs are an example of an Application layer program. In the context of electronic financial markets, the Financial Information Exchange ("FIX") protocol is an Application layer protocol (as well as a session and presentation layer protocol) which may be used with financial systems and financial trading software applications. The FIX protocol comprises a specification for standardized and streamlined real-time exchange of information related to securities transactions and markets.

[0012] Financial instrument trading systems are one example of transaction processing systems that utilize message handling protocols such as TCP/IP and Ethernet, e.g. to receive orders from participants and communicate market data thereto. Generally, a financial instrument trading system, such as a futures exchange, referred to herein also as an "Exchange", such as the Chicago Mercantile Exchange Inc. (CME), provides a contract market where financial instruments, for example futures, options on futures and spread contracts, are traded among market participants, e.g. traders, brokers, etc. Financial instruments, and their corresponding components, are communicated through the CME network as a series of messages. Futures is a term used to designate all contracts for the purchase or sale of financial instruments or physical commodities for future delivery or cash settlement, and which are traded on a commodity futures exchange. A futures contract is a standardized legally binding agreement to buy (long) or sell (short) a commodity or financial instrument at a specified price at a predetermined future time. An option is the right, but not the obligation, to sell (put) or buy (call) the underlying instrument (for example, a futures contract) at a specified price within a specified time. The commodity or instrument to be delivered in fulfillment of the contract, or alternatively the commodity, instrument or reference for which the cash market price shall determine the final settlement price of the futures contract, is known as the contract's "underlying" reference, instrument or commodity, also referred to as the "underlier." The

terms and conditions of each futures contract are standardized as to the specification of the contract's underlier, the quality and quantity of such underlier, delivery date, and means of contract settlement, i.e. physical delivery or cash settlement. Cash Settlement is a method of settling a futures contract whereby the parties effect final settlement when the contract expires by paying/receiving the pecuniary loss/gain of the contract, e.g. by comparing the contract price to the market price or other reference price of the underlier at the time of settlement, related to the contract in cash, rather than by effecting physical delivery, i.e. the actual exchange of the underlying reference or commodity at a price determined by the futures contract.

[0013] Typically, the Exchange provides for centralized "clearing" by which all trades are confirmed and matched, and open positions are settled each day until expired (such as in the case of an option), offset or delivered. The clearing function includes messaging among participants in the market. Matching, which is a function typically performed by the Exchange, is a process, for a given order which specifies a desire to buy or sell a quantity of a particular instrument at a particular price, of seeking/identifying one or more wholly or partially, with respect to quantity, satisfying counter orders thereto, e.g. a sell counter to an order to buy, or vice versa, for the same instrument at the same, or sometimes better, price (but not necessarily the same quantity), which are then paired for execution to complete a trade between the respective market participants (via the Exchange) and at least partially satisfy the desired quantity of one or both of the order and/or the counter order, with any residual unsatisfied quantity left to await another suitable counter order, referred to as "resting."

[0014] A "Clearing House," which is typically an adjunct to the Exchange and may be an operating division thereof, is responsible for settling trading accounts, clearing trades, collecting and maintaining performance bond funds, regulating delivery, and reporting trading data to market regulators and to the market participants. An essential role of the clearing house is to mitigate credit risk via the clearing process. Clearing is the procedure through which the Clearing House becomes buyer to each seller of a futures contract, and seller to each buyer, also referred to as a "novation," and assumes responsibility for protecting buyers and sellers from financial loss due to breach of contract, by assuring performance on each contract. A clearing member is a firm qualified to clear trades through the Clearing House.

[0015] Current financial instrument trading systems allow traders to submit orders and receive confirmations, market data, and other information electronically via a communications network. These "electronic" marketplaces, implemented by, and also referred to as, "electronic trading systems," are an alternative trading forum to pit based trading systems whereby the traders, or their representatives, all physically stand in a designated location, i.e. a trading pit, and trade with each other via oral and visual/hand based communication.

[0016] In particular, electronic trading of financial instruments, such as futures contracts, is conducted by market participants sending orders, such as to buy or sell one or more futures contracts, in electronic form to the Exchange, e.g. using messages comprising the orders. These electronically submitted orders to buy and sell are then matched, if possible, by the Exchange, i.e. by the Exchange's matching engine, to execute a trade. Outstanding (unmatched, wholly

unsatisfied/unfilled or partially satisfied/filled) orders are maintained in one or more data structures or databases referred to as "order books," such orders being referred to as "resting," and generally made visible, i.e., their availability for trading may be advertised, to the market participants through electronic notifications/broadcasts, referred to as market data feeds, also conveyed via messages. An order book is typically maintained for each product, e.g. instrument, traded on the electronic trading system and generally defines or otherwise represents the state of the market for that product, i.e. the current prices at which the market participants are willing buy or sell that product as well as the quantities thereof. As such, as used herein, an order book for a product may also be referred to as a market for that product.

[0017] A market data feed, referred to as market data or market feed, is a compressed or uncompressed real time (with respect to market events), or substantial approximation thereof, data/message stream, i.e. sequence or series of messages, provided by the Exchange directly, or via a third party intermediary. A market data feed may be comprised of individual messages, each comprising one or more packets or datagrams, and may carry, for example, pricing or other information regarding orders placed, traded instruments and other market information, such as summary values and statistical values, or combinations thereof, and may be transmitted, e.g. multi-casted, to the market participants using standardized protocols, such as TCP or UDP over Ethernet. The messages of the market data feed, comprising data relating to an event, may be generated and transmitted responsive to the occurrence of that event, i.e. in real time with respect thereto, or generated and/or transmitted at a later time subsequent to the event and may comprise only the data related to the event or include additional data. More than one market data feed, each, for example, carrying different information, may be provided. Standard protocols that may be utilized for the communication of orders and/or transmission of market data feeds is the Financial Information Exchange (FIX) protocol, and the Financial Information Exchange Adapted for Streaming (FAST), aka FIX/FAST, which is used by multiple exchanges to distribute their market data. Other protocols, standard or proprietary, may also be used for order communication and/or market data feeds in addition to or in lieu of FIX or FIX/FAST, such as Simple Binary Encoding, ITCH/OUCH, etc. Pricing information conveyed by the market data feed may include the prices, or changes thereto, of resting orders, prices at which particular orders were recently traded, or other information representative of the state of the market or changes therein. Separate, directed/private, messages may also be transmitted, in real time or otherwise, directly to market participants to confirm receipt of orders, cancellation of orders and otherwise provide acknowledgment or notification of matching and other events relevant, or otherwise privy, only to the particular market participant.

[0018] As may be perceived/experienced by the market participants from outside the Exchange or electronic trading system operated thereby, the following sequence describes how, at least in part, information may be propagated in such a system, through a series of messages, and how orders may be processed:

[0019] (1) An opportunity is created at a matching engine of the Exchange, such as by placing a recently received but unmatched order on the order book to rest

indicating a trader willing to trade a particular quantity at a particular price, i.e. an opportunity for another trader to trade therewith;

[0020] (2) The matching engine creates an update reflecting the opportunity and sends it to a feed engine;

[0021] (3) The feed engine transmits, e.g. using multicasting, or multiple unicast streams, or other mode of broadcasting, the update in a message sent to all of the market participants to advertise the opportunity to trade, i.e. transmits a broadcast market data feed message;

[0022] (4) The market participants evaluate the opportunity and each, upon completion of their evaluation, may or may not choose to respond with an order, i.e. a message containing their counter-order, responsive to the resting order, i.e. counter to the resting order for some or all of the available quantity;

[0023] (5) The Exchange gateway receives any counter orders generated by the market participants, ascribes an ordering to each received order relative to the other received orders, sends confirmation of receipt back directly to each submitting market participant, and forwards the received orders to the matching engine; and

[0024] (6) The matching engine evaluates the received orders and matches, using for example, a price/time priority algorithm, the first suitable arriving order, based on the ascribed ordering, against the resting opportunity and a trade is executed. Depending upon the available quantity of the resting order and the quantity of the incoming order, one of the resting order or the incoming order may be fully satisfied with any remainder of either order being placed on the order book to rest, or cancelled, etc.

[0025] To attract and retain market participants and encourage participation, electronic trading systems ideally attempt to offer a more efficient, fair and balanced market where market prices reflect a true consensus of the value of traded products among the market participants, and which minimize, if not eliminate, surreptitious or overt subversion, influence of, or manipulation by, any one or more market participants or other entities, intentional or otherwise, and unfair or inequitable advantages, with respect to access to information or opportunities. To accomplish these goals, for example, electronic trading systems may operate in a deterministic, i.e. a causal, predictable, or otherwise expected, manner, as understood and experienced by the market participants, i.e. the customers of the Exchange.

[0026] Accordingly, the operators of electronic trading systems, alone or in conjunction with, or at the direction of, regulatory or industry organizations, typically publish or otherwise promulgate rules or regulations, referred to as business or operating rules, as well as specify protocols for interaction therewith, which govern or otherwise specify the operation of the system. These rules define how, for example, multiple transactions are processed by the system where those transactions have relationships or dependencies there between which may affect the result of such processing. Such business rules may include, for example, order allocation rules, i.e. rules which dictate which of multiple competing resting orders will be matched with a particular incoming order counter thereto having insufficient quantity to fill all of the suitable resting orders. For example, under a first-in-first-out methodology, the first order, of the com-

peting resting orders at the same price, that was received by the electronic trading system will be matched with the incoming counter-order and filled to the extent possible by the available quantity, with any residual quantity of the incoming counter order then being allocated to the next received suitable competing resting order and so on until the available quantity of the incoming counter order is exhausted. However, additional or alternative matching/allocation rules may be implemented as well, for example to ensure fair and equal access, improve trading opportunities, etc., by allocating, such as proportionally, the available quantity of the incoming counter order among all, or a subset, of the competing resting orders until the available quantity is exhausted.

[0027] These business rules may further include specifications of messaging protocols, as described above, used by the electronic trading system for communication with market participants. In particular, these rules may specify which protocols are being used and how they are implemented so that market participants not only understand how to communicate with the exchange, i.e. how to properly format and transmit orders, but also understand how the electronic trading system will respond to their messages as well as process their messages with respect to messages received from other market participants, e.g. such as when messages comprising competing orders are received substantially simultaneously. While the electronic trading system typically may not dictate, specify or control how market participants implement their trading or communications systems, the electronic trading system's own implementation of messaging protocols often defacto-dictates broadly, if not specifically, how a market participant must implement their own messaging protocols in order to successfully interact therewith.

[0028] Once such business rules are established, or modified, market participants will expect, and overseeing regulatory entities may require, that the electronic trading system operate in accordance therewith. That is, if the Exchange adopts a rule to give first arriving orders priority over later arriving orders, a market participant who submits an earlier arriving order will expect their order to be filled prior to a later arriving order submitted by another market participant. It will be appreciated that these rules, by which operators of an electronic trading system may choose to operate their system, may vary at the discretion of the operators, subject to regulatory concerns. Generally, the term "transactional determinism" may refer to the processing, or the appearance thereof, of orders in accordance with the defined business rules.

[0029] In addition to efficiency, fairness and equity, electronic trading systems further provide significant performance improvements allowing for rapid high volume transaction processing which benefits both the Exchange and market participants. Metrics of electronic trading system performance include latency and throughput. Latency may be measured as the response time of the Exchange. This can be measured in a number of different contexts: the time elapsed from when an order, or order cancellation, is received to when a confirmation/acknowledgment of receipt is transmitted, from when an order is received to when an execution notification is transmitted, or the time elapsed from when an order is received to information about that order being disseminated in the market data feed. Throughput may be measured as the maximum number of orders or trades per

second that the electronic trading system can support, i.e. receive and acknowledge, receive and match, etc.

[0030] Generally, market participants desire timely and deterministic market data updates, low latency/high throughput order processing, and prompt confirmations of their instructions to allow them to: competitively, frequently and confidently evaluate, and react to, or, conversely avoid, discrete, finite, fast moving/changing or ephemeral market events; and/or otherwise coordinate, or synchronize their trading activities with other related business concerns or activities, with less uncertainty with respect to their order status. Higher volume capacity and transaction processing performance provides these benefits as well as, without detrimentally affecting that capacity or performance: further improves market access and market liquidity, such as by allowing for participation by more market participants; the provision of additional financial products, and/or additional markets therefore, to meet the varying needs of the market participants; and efficient identification of additional explicit and implicit intra- and inter-market trading opportunities. The Exchange benefits, for example, from the increased transaction volume from which revenue is derived, e.g. via transaction fees.

[0031] Current electronic trading systems already offer significant performance advantages. However, increasing transaction volumes from an increasing number of market participants, implementation by some market participants of algorithmic and/or electronic liquidity provisioning methodologies whereby high speed computers automatically monitor markets and react, with deterministic efficiency, i.e. minimal latency, to events, coupled with a continued demand for ever-decreasing processing latencies and response times, is driving a need for additional capacity and performance improvements to maintain performance as experienced by each market participant and avoid detrimental consequences, such as capacity exhaustion and inequitable access. For example, the increasing efficiency with which market participants may evaluate and respond to changes in market data, such as responsive to a market event, is increasing the rate at which transactions are received by the electronic trading system, narrowing the delay of receipt gap there between and necessitating a need for a higher degree of discrimination so as to resolve the order/sequence in which those transactions are received, upon which the deterministic operation of the electronic trading system may be based, e.g. for order priority, allocation, etc. Furthermore, the addition, by electronic trading systems, of additional communication bandwidth, in an effort to equitably increase capacity and opportunity, allows for more transactions to be submitted over multiple parallel paths into the system which may then all converge at a common point of determinism. Accordingly, not only must the electronic trading system discriminate among closely received incoming transactions, but must further arbitrate among a congestion of transactions received simultaneously, or temporally so close together as to be considered simultaneously received, i.e. the difference in their time of receipt being too close to measure by the implemented discrimination mechanisms, also referred to as “substantially simultaneously”.

[0032] In addition to increased capacity and lower latency, the global nature of business has further driven a need for fault tolerance to increase availability and reliability of

electronic trading systems. Scheduled outages must be minimized and unscheduled outages must be eliminated.

[0033] Furthermore, to implement the Exchange’s clearing function, which mitigates the concerns of market participants relating to performance by counter parties, protects the interests of the Exchange and otherwise adequately manages/mitigates risk, risk management systems having corresponding operational efficiency and performance are needed so as to protect the Exchange from loss while minimizing impediments to market operations or distractions to market participants with ancillary and unnecessary tasks. In addition, increased transaction volume may further translate into greater exposure for market participants requiring greater amounts of capital to be posted to cover losses.

[0034] Without additional safeguards, improved speed and efficiency may, in some circumstances, improve the speed at which problems may occur, grow and propagate, or otherwise may be exploited, such as where the market ceases to operate as intended, i.e. the market may no longer reflect a true consensus of the value of traded products among the market participants. Such problems are typically, but not always, evidenced by extreme market activity such as large changes in price, whether up or down, over a short period of time or an extreme volume of trades taking place. In particular, market participants, whether human or electronic, may not always react in a rational manner, such as when presented with imperfect information, when acting in fraudulent or otherwise unethical manner, and/or due to faulty training or design.

[0035] For example, while communications technologies may have improved, inequities still exist in both access to information and access to opportunities to participate, which may not be due to any violations of legislative, regulatory and/or ethical rules, e.g. some traders receive information before other traders because they have access to, such as due to physical location, or otherwise can afford faster communications channels, some traders may be able to place trade orders more quickly than others because they have faster computers, etc. In many cases, certain trader behaviors may be triggered by a market event, such as a change in price, creating a feedback loop where an initial reaction may then cause further market events, such as continued price drops, triggering further responses and resulting in an extreme change in the price of the traded product in a short period of time. High speed trading may exacerbate the problem as there may be little time for traders/algorithmic trading systems, or those overseeing them, to contemplate and temper their reactions before significant losses may be incurred. Furthermore, improved communications among traders facilitates a continuous stream of information and propagation of behavior, efficiency seeking or not, in one market to other markets as traders in those other markets react to the results of that behavior.

[0036] Market protection systems may therefore be needed to monitor and evaluate trading activity, detect undesired activity and appropriately react more quickly to mitigate the spread of problems, again without impeding legitimate market operation.

BRIEF DESCRIPTION OF THE DRAWINGS

[0037] FIG. 1 depicts a conceptual model for the lifecycle of a data packet, or collection of data packets, traveling from one computer to another computer.

[0038] FIG. 2 depicts an illustrative environment in which the system that may be used to implement aspects of the disclosed embodiments is located.

[0039] FIG. 3A depicts an illustrative electronic trading system with which the disclosed embodiments may be utilized.

[0040] FIG. 3B depicts an alternative illustrative electronic trading system with which the disclosed embodiments may be utilized.

[0041] FIG. 4A depicts an illustrative implementation according to the disclosed embodiments.

[0042] FIG. 4B depicts a functional implementation of the system of FIG. 4A.

[0043] FIG. 5 an illustrative set of instructions for use with the system of FIGS. 1-4 according to one embodiment.

[0044] FIG. 6 depicts an illustrative buffer for use with the system of FIGS. 1-5 according to one embodiment.

[0045] FIG. 7 depicts a flow chart showing operation of the system of FIGS. 1-6 according to one embodiment.

[0046] FIG. 8 shows an illustrative embodiment of a general computer system for use with the system of FIGS. 1-7.

[0047] FIG. 9 depicts example scenarios to which the disclosed embodiments may be applied.

DETAILED DESCRIPTION

[0048] The disclosed embodiments relate to implementation, such as by a message processing system architecture, of a message handling system and/or protocol which mitigates optimistic messaging behavior. As used herein optimistic messaging behavior may at least refer to the transmission, in whole or in part, of a message, or of one or more messages of a sequence thereof, to a receiving system, such as the system implemented by the disclosed embodiments, prior to the sender determining that the message(s) is/are desired, alone or in conjunction with subsequently canceling or otherwise invalidating the message(s) once it has been determined that it is not desired. The disclosed embodiments mitigate such behavior by detecting and taking action with respect to these types of messages to deter optimizing behavior which may or may not be manipulative, while maintaining reliable message handling under increasing processing loads with minimal impact on users being able to send messages to the system.

[0049] The disclosed technology addresses the need in the art for mitigating optimistic messaging which is unique to electronic communications systems which decompose messages into transmissible portions or otherwise transmit messages in sequence and use protocols to ensure reliable, complete and/or error-free transmission thereof, as opposed to prior trading systems which relied upon traders communicating face to face or via telephone where messages were communicated there between in their entirety. Specifically, the disclosed technology solves a problem that uniquely arises in the fields of computer technology, electronic communications networks and exchange computing systems, where data transmission and communications protocols may be managed or optimized to alter or otherwise control message ordering and/or delivery, in particular with respect to messages communicated to same destination from other sources. Thus, the disclosed solution is rooted in computer technology in order to overcome a problem specifically arising in the computer systems and electronic communications networks used by such electronic trading systems.

Indeed, the subject technology improves the functioning of the computer by, for example, detecting message optimization and segregating such optimization from incidental or accidental communications artifacts, allowing the computer to use fewer resources.

[0050] The disclosed embodiments are drawn to systems and methods that include specific computing components, each being specially programmed to perform a technological function as part of a greater technological process. The disclosed embodiments include separate system components interconnected in a specific way to implement aspects of the disclosed system and include sufficient specific structure and function and, as such, are not drawn to an abstract idea.

[0051] The disclosed embodiments are not directed to any method for “obtaining, transforming and determining,” which is involved in all computing functionality. The disclosed embodiments and features recited in this regard provide numerous advantages. The instant embodiments do not preempt all methods of “obtaining, transforming, and determining,” and are specifically directed towards the disclosed functionality. The disclosed embodiments implement specific rules and features that improve the operation of a particular genus of a technological process, which does not preempt all techniques of obtaining, transforming and determining, which, at some level, is part of every computing process.

[0052] The disclosed messaging system architecture, as will be described, may be implemented using computer technology, e.g. processors, memories, electronic communications networks, inter or intra device, e.g. node, communications, and the like. As used herein, the terms “micro-processor” or “general-purpose processor” (“GPP”) may refer to a hardware device that fetches instructions and data from a memory or storage device and executes those instructions (for example, an Intel Xeon processor or an AMD Opteron processor) to then, for example, process the data in accordance therewith. The term “reconfigurable logic” may refer to any logic technology whose form and function can be significantly altered (i.e., reconfigured) in the field post-manufacture as opposed to a microprocessor, whose function can change post-manufacture, e.g. via computer executable software code, but whose form, e.g. the arrangement/layout and interconnection of logical structures, is fixed at manufacture. The term “software” will refer to data processing functionality that is deployed on a GPP. The term “firmware” will refer to data processing functionality that is deployed on reconfigurable logic. One example of a reconfigurable logic is a field programmable gate array (“FPGA”) which is a reconfigurable integrated circuit. An FPGA may contain programmable logic components called “logic blocks”, and a hierarchy of reconfigurable interconnects that allow the blocks to be “wired together”—somewhat like many (changeable) logic gates that can be inter-wired in (many) different configurations. Logic blocks may be configured to perform complex combinatorial functions, or merely simple logic gates like AND, OR, NOT and XOR. An FPGA may further include memory elements, which may be simple flip-flops or more complete blocks of memory

[0053] Messages handled by the disclosed architecture may be related to business transactions which, themselves, may be defined as one or more operations or acts which are undertaken according to one or more associated business rules (including industry, legal or regulatory requirements or customs) to accomplish a business or commercial purpose,

which may include compliance with industry, regulatory or legal requirements. A business transaction may be implemented by one or more computer processing and/or database operations/program steps, which themselves may be referred to as transactions. Business transactions, as defined by the associated business rules, may be characterized as deterministic in that they be characterized by an interdependency or relationship which affects their result, such as a dependency on the order in which they are processed, such as a temporal order, and/or a dependency on real time processing, as defined by the business rules, so as to effect the business/commercial purpose and/or meet participant expectations, referred to herein as “transactional determinism.” Generally, a set of deterministic transactions may provide a particular result when executed in one order and a different result when executed in a different order. In some applications, deterministic processing may be preferred/prioritized over real time processing. For more detail on deterministic operation in a trading system, see U.S. patent application Ser. No. 14/074,675, filed on Nov. 7, 2013, published as U.S. Patent Publication No. 2015/0127516, entitled “Transactionally Deterministic High Speed Financial Exchange Having Improved, Efficiency, Communication, Customization, Performance, Access, Trading Opportunities, Credit Controls, And Fault Tolerance”, the entirety of which is incorporated by reference herein and relied upon.

[0054] For example, wherein the business rules define a first-in-first-out (“FIFO”) process for matching offers with counter-offers to effect an exchange or trade, when an offer transaction is received to which no prior counter offer transaction has been previously received, it should be matched with the next suitable counter-offer transaction received rather than a later received counter-offer transaction. It will be appreciated that the processing of a given transaction may involve delaying further processing of that transaction in favor of a later received transaction, such as dependent upon conditions specified by the earlier transaction and/or the defined business rules. It will further be appreciated that, at a minimum, any representation of the current state of a business environment, e.g. market, or changes thereto, in which the business transactions are transacted should present an accurate reflection of the actual state or state change in accordance with the defined business rules. In the disclosed embodiments, the phrase “financial transaction” will refer to a business transaction involving the purchase or sale of financial instruments, such as futures contracts or options thereon, spread or other combination contracts and the like, as will be described. As was described above, electronic trading systems generally define their business rules and then must ensure transactional determinism in compliance therewith.

[0055] Alternatively, as will be described, the embodiments described below may be implemented using FPGA’s or other reconfigurable logic. Implementing processing tasks and algorithms using an FPGA can yield significant performance enhancements over implementations using traditional microprocessors and operating systems. In particular, an FPGA based system implementation may avoid the processing overhead and uncontrollable/unnecessary optimizations implemented by general purpose processors, compilers, operating systems and communications protocols, as well as the security vulnerabilities thereof. For example, an FPGA may avoid interrupt handling, error correction, pre-fetching and other unnecessary microprocessor operations/optimiza-

tions, as well as generic processing/housekeeping tasks of the operating system which are not needed, such as garbage collection, unnecessary memory swaps, cache loads, task switching, cycle stealing, etc. Further an FPGA implementation may avoid the use of general purpose compilers which may introduce, for example, undesired program code optimizations.

[0056] For example, using an FPGA based implementation may permit components of a trading system to be collocated, such as via a custom interface, or otherwise closely interconnected with networking equipment, such as a router or switch, e.g. via a backplane thereof. This would allow the trading system to have access to incoming transactions as quickly as possible and avoid the latency introduced, not only by having to route the transaction over conventional networking media, but also by the communications protocols, e.g. TCP, used to perform that routing. One exemplary implementation is referred to as a “Smart Network Interface Controller” or “SmartNIC” which is a device which typically brings together high-speed network interfaces, a PCIe host interface, memory and an FPGA. The FPGA implements the NIC controller, acting as the bridge between the host computer and the network at the “physical layer” and allows user-designed custom processing logic to be integrated directly into the data path. This may allow a smart NIC to function as a programmable trading platform under the supervision of a host CPU. Under the OSI model, which, as described above, is a conceptual model that characterizes and standardizes the internal functions of a communication system by partitioning it into abstraction layers, the physical abstraction layer defines electrical and physical specifications for devices. In particular, it defines the relationship between a device and a transmission medium, such as a copper or fiber optical cable. This includes the layout of pins, voltages, line impedance, cable specifications, signal timing, hubs, repeaters, network adapters, host bus adapters (HBA used in storage area networks) and more. The major functions and services performed by the physical layer include: establishment and termination of a connection to a communications medium; participation in the process whereby the communication resources are effectively shared among multiple users, for example, contention resolution and flow control; and modulation or conversion between the representation of digital data in user equipment and the corresponding signals transmitted over a communications channel, these signals operating over the physical cabling (such as copper and optical fiber) or over a radio link.

[0057] An electronic trading system, as described above, is an example of a transaction processing system which may employ the disclosed message processing architecture. For example, in normal operation of an electronic trading system, as will be described with reference to FIGS. 1-3B, such as may be implemented by the Exchange described above, the message handling system and messaging protocols implemented thereby may function to allow users, e.g. market participants or traders, to send messages to the Exchange comprising orders to trade, i.e. orders to buy or sell quantities of one or more tradeable items, such as financial instruments, at specified prices, requests to cancel previously sent orders, as well as other communications. The source, or sending host 102 as depicted in FIG. 1 may be a market participant 102a . . . n shown FIGS. 2-4B such as a trader, i.e. the trading system/computer operated

thereby, while the destination, or receiving host **124** of FIG. **1** may be an Exchange/electronic trading system **208**, or the receiver **206** thereof as shown in FIGS. **2-4B**, i.e. the devices utilized by the Exchange/electronic trading system **208** to receive and process messages received from participant computer systems **102**.

[0058] In the example operation, a user/trader who decides that they wish to execute a trade order inputs, calculates or otherwise directs the creation of the necessary information **126** to complete the trade order, e.g. by using a trading software application, implemented at the Application Layer **104**, executing on their source/computer system **102/102a-n** which is connected to the Exchange **208**/Receiving Host **124**, such as via, as shown in FIGS. **2-3B**, an interconnecting network infrastructure **134** which may include the private and/or public network infrastructure **204a** between the sources **102a-n** and the electronic trading system **208**, i.e. external thereto, as well as any of the internal network infrastructure **302** of the electronic trading system **208** which couples the external infrastructure with the receiving host **124**/receiver **206**. This may be the network **820** described in more detail below with respect to FIG. **8**. This information **126** may include, for example, data indicative of the instrument they wish to transact, the side (buy or sell), the price, and the quantity. As will be described, a user/trader, or their computer system **102a-n** may create an order to trade based on information available to them regarding an event, such as the state of a marketplace, e.g. based on data received via a market data feed and/or news feed, and/or other information such as based on the result of a calculation. The trading software application may be configured to format the order/message according to the FIX protocol described above. This may include calculating and including an error detection and/or correction code, such as a checksum value, and including it as part of order/message as dictated by the application protocol such that the application, e.g. the trading software application, may validate the order/message upon receipt.

[0059] The user's computer system **102** may implement a protocol stack conforming to the OSI model, as was described above, and depicted in FIG. **1**, executing software components which implement the various layers, or a subset thereof, of the model, e.g., the layers **104**, **106**, **108**, **110**, **112** depicted in FIG. **1**. The order is input into, or otherwise computed by the trading application which, in this example, is implemented at the application layer **104** of the OSI model, which then generates the information **126**, i.e. a FIX formatted message, including the requisite checksum value, for transmission to the Exchange.

[0060] The information **126** comprising the order, e.g. the FIX message, is conveyed from the application layer **104** to the transport layer **106** of the user's computer system **102** protocol stack via the interaction between the trading application and the software component(s) which implements the transport layer **106**. Upon receipt, the information **126** is reorganized, or packaged, by the software component implementing the transport layer **106** according to the implemented transport protocol, e.g. TCP. This process is typically referred to as data encapsulation and includes segmenting the message, as will be described below, into one or more portions depending upon the message size and augmenting each of portions of the message with additional data such as the source of the message and the destination for the message. In one example, the transport protocol is TCP. In

another example, the transport protocol may be UDP. With respect to TCP, the information **126** is converted into one or more TCP segments, also more commonly known as TCP packets, such as the TCP segment **128** of FIG. **1**.

[0061] The amount of data that may be sent in each segment/packet may be limited. Accordingly, based on the size of the information **126** to be transmitted, e.g. the size of the FIX message, the software component implementing the transport layer **106** on the user's computer **102** determines how many packets will be required to send the information **126**. It will be appreciated that segment/packet size may in fact exceed the size of the information **126**, e.g. the size of a message/order, such that multiple messages/orders may be included in a single packet. Furthermore, to maximize packet utilization where the packet size exceeds the size of one or more messages/orders with excess capacity less than an additional whole order/message, the additional order/message may be apportioned across multiple packets along with the whole orders/messages included therein. According to TCP, the information to transmit may be combined into one, and/or broken up into multiple, packets according to the size of the information and the maximum transmission unit size of the data link layer of the network to which the sender is connected. The transport layer **106** interacts or otherwise negotiates with the data link layer **110** to determine what maximum transmission size unit to use. According to the protocol specification, the largest size unit allowed by TCP is 65535 bytes, however, in practice the size of a packet may be much smaller because protocols on the network link layer may not support packets that are 65535 bytes. The typical size for TCP packets may be between, at a minimum, twenty bytes (which is the size of a TCP header) and 1500 bytes. After the transport layer **106** determines the maximum transmission unit size, the information **126** to be transmitted is broken up into packets that do not exceed the maximum transmission unit size. For example, if the information is to be sent over Ethernet, then the maximum size for any given TCP packet cannot be larger than the size of the maximum transmission unit size for the Ethernet protocol, which is 1500 bytes. Once the information is divided into packets, each packet is assigned a sequence number, i.e., in accordance with the sequence number agreed upon between the sender and receiver in the three way handshake described below.

[0062] The transport layer **106** of the user's computer system **102** further begins the process of establishing a connection, referred to as a "session," with the destination computer **124**. Establishing a connection includes exchanging messages between the source **102** and destination **124** computers to ensure that both computers are ready to exchange information. This exchange involves a request sent from the source **102** to the destination **124** computer to establish a logical connection, a reply from the destination **124** computer and an initial sequence number (chosen at random) to be used as part of the data to be transmitted, and an acknowledgement from the source **102** computer of the receipt of the reply and the initial sequence number to use for transferring data in the session to the destination **124** computer. This exchange of messages between source **102** and destination **124** computers is referred to as a "three way handshake." The established connection opens a session between the source **102** and destination **124** computer during which information may be exchanged back and forth in what is known as full-duplex communication. Once the connec-

tion is established, transmission of the user's information, in this case the market order data **126**, which has been packaged into at least one or more TCP packets **128**, to the destination computer **124** commences as will be described.

[0063] Once segmented into one or more packets, each TCP packet **128** is conveyed to the software component which implements the Network layer **108** of the sending host **102**. The Network layer **108** may, for example, implement the Internet Protocol ("IP"), such as the IPv4 or IPv6 variant. The IP protocol prepares each of one or more the TCP packets **128** for delivery by formatting the TCP packet **128** into an IP datagram **130**. The preparation consists of augmenting the TCP packet **128** with transportation information to form an IP datagram **130**, such as what ports to use for communication, an IP header, and other information about the IP datagram such as sending and receiving hosts. Once each IP datagram **130** is properly formatted it is conveyed to the software component which implements the data link layer **110** for the sending host **102**.

[0064] At the data link layer **110**, the IP datagram **130** is handled according to the particular networking protocol used by the communication network **134** which, in one example, may be the Ethernet protocol. According to the Ethernet protocol, the IP datagram **130** is transformed into an Ethernet frame **132**. As part of this process, the IP datagram **130** is augmented with a frame header and a frame footer. The frame header includes addressing information as may further include data specifying the type of packet (EtherType) and/or the amount of data being conveyed by the frame (with Ethernet, the amount of data conveyed by a given frame is implicit in the signaling protocol via the receipt of end-of-frame/FCS indicator). The frame footer includes a Frame Check Sequence ("FCS"), generated via a cyclic redundancy check ("CRC") algorithm, which is used by the destination computer **124**, or intermediate networking devices/nodes, to check if the message was corrupted in transit from source to destination. Using the information that is included in the Ethernet frame **132**, at least part of the user's original market order **126** is now ready to be physically transmitted to the destination computer **124**, or receiving host **124**.

[0065] Each Ethernet frame **132**, which may carry one of the one or more TCP packets **128** or FIX messages, travels over a physical network layer **112** via a network connection/media **134**, e.g. via the packet switched network described above, until it eventually reaches the physical network layer **114** of the receiving destination **124**. That is, each frame **132** is transmitted via the physical network port of the sending computer **102** over the physical wired and/or wireless networking medium coupled therewith and to the physical network port of the destination system **124**. En route, the frame **132** may pass through various switches, gateway or router devices and various physical media types. These devices are generally coupled between multiple network paths, i.e. have numerous inputs coupled with different upstream paths and numerous outputs, each coupled with a different downstream path. The devices determine the destination of the various data received from the different inputs and route that data to the proper output coupled with the network path leading to the destination. As such, these devices cause frames/packets having a common destination received from multiple inputs to converge on a common output coupled with that destination.

[0066] Ultimately, the frame is received at the destination **124**, e.g. the Exchange **208**, such by an ingress switch device **206** or other device. The receiving host's **124** protocol stack reverses the above described process in order to provide a match engine **308**, coupled therewith, with the user's trade order **126**.

[0067] In particular, as each Ethernet frame **132** is received by the hardware and/or software component which implements the physical and data link layers **114** and **116** respectively of the destination host **124**, the FCS for the Ethernet frame **132** is computed. Then the FCS included in the received Ethernet frame **132** is checked at the data link layer **116** against the computed FCS to ensure that the frame **132** was not corrupted in transit. The FCS may be checked via a running process which computes the FCS as the frame **132** is received and results in a specific value if the FCS is correct. If the received and computed FCS values differ, the frame **132** will typically be discarded. As will be described, in a typical implementation, it is the responsibility of the higher layers of the protocol stack, e.g. the TCP layer **118**, to cause retransmission of the lost frame **132**. If the frame **132** is uncorrupted, the Ethernet software program on the receiving computer **124** removes the Ethernet header and footer from the Ethernet frame **132**, and the remaining IP datagram is conveyed to the Network layer **118** of the receiving computer **124**.

[0068] When multiple frames **132** directed to the same destination are received, as was described above, substantially simultaneously via different inputs of a receiving intermediary device, such as a gateway or network switch, operative to direct the frames **132** to/toward the destination device, that intermediary device may be required to arbitrate among those received frames **132** to determine which to send first to its output coupled with the destination device or the path thereto. This arbitration is typically programmed into the intermediary device and may comprise a round-robin, pseudo random or other algorithm for determining which frame **132** is transmitted/forwarded to the destination first.

[0069] At the Network layer **118**, the IP datagram **130** is received. The IP header is read to determine the source IP address and destination IP address as well as other data needed to route the IP datagram **130**. Prior to sending the information to the transport layer **120**, the IP software program **118** (or hardware based IP or TCP/IP processor) on the destination computer **124** will remove the IP header from the IP datagram **130** to extract the payload, e.g. the TCP packet, carried thereby which is then conveyed to the transport layer **120**.

[0070] The software/hardware which implements the transport layer **120** of the destination computer **124** provides/implements buffers which are used to store packets as they are received during the TCP session. As the transport layer **120** of the destination computer **124** receives packets, the software program implementing the transport layer **120** of the destination computer **124** sends acknowledgment messages to the source **102** acknowledging that each packet has been received. The acknowledgments are cumulative, which means, for example, that if the receiver sends an acknowledgment that a packet with sequence number ten has been received, this indicates also that packets with a sequence number less than ten have also been received and that they do not need to be sent or resent. If a packet is missing, e.g. a later sequenced packet is received before an

intervening sequenced packet, an acknowledgment will not be sent causing the sender **102** to eventually resend the unacknowledged packets or all of the packets depending upon the configuration, as was described above. The received packets are put into a TCP buffer, segregated by session, where they may end up being reorganized into the correct sequence order. The receiving host **124** may know that no more packets are forthcoming from the source **102** when the source **102** indicates no more information, or packets are to be sent but such an indication may not always be received. The received packets are put in sequence order that corresponds to the initial sequence number that was decided in the three way handshake and the sequence numbers included in each packet. At the transport layer **120** the header information is read to determine which application at the application layer **122** is to receive the information. The software program implementing the transport layer **120** of the destination computer **124** removes the header information for each packet that is received by the receiving host, and puts all of the received packets in sequence order. Once the received packets are in sequence order, after their headers are removed, the bytes that make up the received packets are streamed to the correct application at the application layer **122** on the receiving host. The application on the application layer **122** of the receiving host receives the stream of bytes from the transport layer **120** and interprets the stream of bytes in a way that is meaningful to the application program. In particular, the application may determine whether it has received a complete message, e.g. by looking for and validating a checksum value. For example, a trading application may determine whether a proper and complete FIX message has been received. As the transport layer **120** will continue to receive packets and provide them to the application, the application may need to continually check to determine whether a complete message has been received before further processing may be commenced.

[0071] In particular, once the complete FIX message has been received, and its checksum validated, the user's order **126** is provided to the order input processing functionality of the Exchange, such as the orderer **310** described in more detail herein, which may, for example, ascribe a time of receipt to the order in accordance with the deterministic processing of the Exchange.

[0072] In some cases, the software program implementing the transport layer **120** on the destination computer **124** may have multiple TCP sessions open, e.g. where it is simultaneously communicating with multiple sources **102**, one session will be opened to handle the receipt of packets from each source **102**. The packets from each source **102** may be stored/accumulated in separate buffers as they are received. For example, the Exchange may be receiving orders from multiple traders at any given time. When multiple frames **132** arrive, there is the chance of a collision occurring. A collision occurs when two frames **132** attempt to access the same medium at the same time. In the event of a collision the source **102** of each frame will attempt to retransmit their frame **132** at a random interval of time. In a situation where the multiple sources **102** are on the same physical network as each other and the destination computer there may not be collisions, because the software running on the data link layer **116** may set up domains for the transmission of frames **132** specific to each source **102**.

[0073] Typically, the various software components which implement the protocol layers described above are implemented according to defined standards so as to interoperate as described, e.g. the operations of the protocol stacks of the sending and receiving devices **102** and **124** respectively operate in conjunction with each other to ensure reliable and efficient message delivery. However, it may be recognized that, when used with a deterministic transaction processing system, such as the electronic trading system described herein, wherein efficient and timely submission of transactions requests, e.g. orders, is desirable so as to capitalize upon available opportunities before others, the operators of sending devices **102** may be motivated to modify the operations of their protocol stack so as to optimize the operation thereof and the corresponding responsive operation of the receiving device's **124** protocol stack, in order to gain an advantage when submitting their orders.

[0074] For example, as shown in FIG. 9, in many cases, a trader may wish to submit one or more orders to trade responsive to the occurrence of an event, such as based on receipt of information, e.g. a price change received via a market data feed, or information regarding a news/current event such as the announcement of an interest rate change, or based on the result of a computations/calculation such as calculation of a price curve, etc. Typically, once a trader, e.g. Sender B, has knowledge of the occurrence of the event, they follow the process described above to create and submit one or more orders to trade based thereon. The event may alternatively be referred to as a "trigger" as the occurrence of the event triggers other activity, e.g. the submission of a trade order or other transaction request. It will be appreciated that the event may be just the initial receipt of the indication of the occurrence of another event, e.g. the initial receipt of the data, i.e. one or of the bits/bytes/signal, of a market data message announcing a change in the market. The event may be one of a series of pre-cursor events leading to a final or cumulative event, each of which may be considered an event herein. Alternatively, or in addition thereto, the event may be considered generation, receipt or perception of the minimal amount of information about which to make an event dependent decision as contemplated herein.

[0075] However, some, e.g. Sender A, may recognize that these trade order messages, e.g. the FIX message, contains numerous pieces of data and, due to the size of the message, e.g. the amount of data to be transmitted, the message will take a certain amount of time to transmit. In this case, the application layer or TCP protocol may be optimized, e.g. manipulated, for all or a subset of the trade order messages, to generate, prior to the occurrence of the event, one or more TCP packets comprising the event independent data, such as the trader identifier, source information, etc., and transmit at least some of those packets first before the event occurs as shown in FIG. 9. Those packets may be constructed to appear, based on the sequencing data, etc., as if they were constructed based on the entire FIX message when, in fact they were not. Once the event occurs and any event dependent order data is determined, such as price, quantity, side, instrument, message checksum, etc., the TCP software is directed to generate one or more TCP packets comprising the event dependent data and transmit those packets to complete the FIX message, the late arriving packets merely appearing to the receiving device as having been delayed in transit. As compared to the overall size of the FIX message,

these remaining TCP packets comprise a smaller amount of data and therefore take less time to transmit. Accordingly, relative to the time of the occurrence of the event, the trader, e.g. Sender A, is sending less post-event data to the Exchange and therefore is able to send that data faster, e.g. having their order considered to be received first before the order of Sender B as shown in FIG. 9. To the receiving device 124, its protocol stack, unconcerned with the occurrence of the event, merely sees the packets of the complete message arriving out of order and perhaps at irregular intervals, which of course, it is designed to handle as described above.

[0076] Alternatively, or in addition thereto, the application layer protocol may be optimized in concert with the TCP protocol. For example, a trader, e.g. Sender B in FIG. 9, may create/optimize a message size or otherwise create multiple orders/messages, each of which requires a checksum validation value to be included therewith, where the messages are sized such that a particular order message will be split across more than one packet, e.g. a portion, or the majority, of the message is included in a first packet, while the remainder, including the checksum value for that message, is included in a subsequent packet. The trader, e.g. Sender B, may send the first packet(s), e.g. containing the majority of the message, but intentionally hold back the remaining packet, or otherwise cause the packet to be held back such as by not providing the data to the TCP layer, having the checksum value until such time as they determine that they do or do not want the order/message. If they do want it, they transmit the proper checksum value as shown in FIG. 9. If they do not want the message, they transmit a corrupted or otherwise invalid checksum value.

[0077] Further, it may also, or alternatively, be recognized that a trade order message, as broken down into packets/frames, etc., as described above, may be corrupted in transit, such as due to faulty intermediate hardware or software, and need to be resent, again circumstances the above described protocols are designed to handle. Further it may be recognized that messages, packets or frames may be intentionally corrupted. For example, prior to occurrence of an event, a trader may cause their sending system 102 to transmit one or more speculatively generated messages, e.g. a message or order generated based on a speculated outcome of the event, whereby at least part of the speculatively generated message is reserved, e.g. held back. For example, a trader may determine that an event has multiple different outcomes and create one or more speculative orders, or one or more TCP connections for the communication thereof, based on each possible outcome. Once the event occurs and the trader knows of the actual outcome, the trader may determine whether or not they actually want any of the speculatively generated messages. For those messages they want, they correctly transmit the reserved portion whereas for any messages they do not want, they do not send or otherwise corrupt the reserved portion so as to cause the receiving protocol stack to invalidate the portion already received. For example, a user may construct an Ethernet frame comprising a payload of data indicative of a speculative order based on trader's guess as to the outcome of the event. They may then begin transmitting this Ethernet frame, reserving the FCS until the outcome of the event is known. Once the event occurs, concludes or otherwise finishes occurring, if the trader still wants the outgoing/in-process order, they transmit the correct FCS for the outgoing packet. However, if

they do not want the order, they can transmit an incorrect FCS causing the receiving device 124 to discard the frame or otherwise cause the packet to not be processed by the electronic trading system. Similarly, the trader may cause the transmission of all but a reserved portion of the TCP packets which constitute a FIX message of their speculative order. Once the event occurs, if the trader still wants the order, they correctly transmit the remaining portion whereby, not only has the previously transmitted portion potentially been received prior to a competing order but the transmission of the remaining portion requires less time to transmit thereby, with respect to the timing of the occurrence of the event, allowing the trader to submit their order more quickly, potentially, for example, with a minimum sized packet with a minimal or the absolute smallest serialization time through the network. However, if they no longer want the order, they may corrupt the TCP packet or otherwise terminate the TCP session with the receiving device 124, causing the receiving device 124 to discard any received data.

[0078] In electronic trading systems which implement TCP in a mode which enforces packet sequencing, a user may transmit multiple TCP packets of a sequence, each containing, for example, a trade order, but intentionally send those TCP packets out of order, e.g. reserve the first TCP packet of the sequence, knowing the receiving system will not process, i.e. will block processing, those TCP packets of the sequence that were received until the missing TCP packet of the sequence is received. For example, a user may create a sequence of five TCP packets each containing a partial, single or multiple trade orders and transmit packets 2-5 to the receiving system, reserving packet 1. The receiving system will receive and buffer, but then block from further processing, packets 2-5 until the user finally sends packet 1. In this way, a user may front-load multiple messages, e.g. multiple trade orders, in advance of an event, such as to secure priority of processing in the receiving system's buffer, and control when those messages are processed, e.g. subsequent to the occurrence of the event, by transmitting the remaining out of order packet. This may be combined with the above described transmission of incomplete messages whereby the trader frontloads messages, some or all of which maybe incomplete, as was described above, e.g. missing event dependent data such as a price or checksum. In this manner, the sender may not only control when a batch of frontloaded messages are processed, via the sending of an out of sequence message, but which of those messages will be processed, by sending the remaining portions, e.g. the missing bits/bytes, to complete only those messages of the batch they want and, for example, corrupting the remainder.

[0079] The scenarios described above may be referred to as "optimistic messaging". In these cases, at least some of the event independent messages or portion of one or more messages is transmitted before the event occurs, i.e. before the sender knows that they want the message. Subsequent to the event, the remaining messages or event dependent portion, and any remaining event independent portion, is transmitted which may complete or corrupt the messages or the portion previously sent depending upon desire/intent of the sender. As will be described, the event independent portion may include any data which has no dependence, or which is not intended to have dependence, upon the occurrence or result of the event or any data which is sent prior to the occurrence of the event which may include speculative data

generated based on an anticipated occurrence or result of the event, whereas the event dependent portion includes any data generated, or intended to be generated, based on the occurrence or result of the event or any data transmitted after the occurrence or result of the event.

[0080] Other examples of such optimistic messaging behavior may be further evidenced by 1) a user of the system continuously streaming Ethernet frames at or near the maximum transmission rate and corrupting the sent Ethernet frames by sending an invalid FCS for the Ethernet frame, or 2) a user beginning to send out a to-be-corrupted Ethernet frame on the receipt of every incoming packet (before its full payload has been received).

[0081] As was described above, the messaging protocols are designed to handle intermittent problems with the communication of messages across networks, e.g. packet loss, bit loss, data corruption, race conditions, multi-path latency variance, etc. Further, it is understood that the mechanisms implemented to handle these problems come at a certain cost/overhead, e.g. additional network congestion and bandwidth consumption due to retransmission, computational processing consumption due to having to continually evaluate and determine when messages are complete or not and generate retransmission requests, etc. However, when the protocols are optimized so as to unnecessarily trigger these remediation mechanisms, resources, particularly of the receiving devices **124**, are unnecessarily wasted which may then degrade the performance of those devices, in particular with respect to their ability to interact with other traders not engaging in such activity.

[0082] More particularly, message protocols, such as those used to implement each layer of the protocol stack, are used to ensure how information is transmitted along a network, in particular the order in which information is transmitted and the way in which information is transmitted. In some instances users of a system may seek to optimize, or otherwise capitalize on, the intricacies and/or operation of a message protocol, such as by utilizing optimistic messaging, i.e. sending a message or starting to send a message before a user actually knows that they want to send the messages, etc., i.e. before the user has the information the user needs in order to know if the user wants the message to be sent, thereby taking advantage of the “reliability” features of the protocol. For example, a user may transmit a portion of a message, but not transmit the remaining portions of the message until the occurrence of an event. Further, a user may intentionally optimize the order in which packets, containing whole messages or portions thereof, of data are sent to a receiving computer **124** so as to reserve the event dependent portions thereof for transmission only once the event has occurred. In doing so, the bulk of a message or messages may be transmitted in advance of an event with a smaller portion being transmitted later, only after the event has occurred, the smaller portion taking less time to transmit and thereby effectively decreasing the latency of transmission of the overall message with respect to the timing of the occurrence of the event. Alternatively, or in addition thereto, the sender, depending upon the outcome of the event, may determine that they no longer wish to send the message at all and, therefore, take some action to cancel the previously transmitted portion, such as by terminating communications with the receiver or intentionally corrupting the remainder of the message to cause the receiver to discard the message.

[0083] For example, in some embodiments the transaction processing system may be a financial exchange, like the Exchange **208** of FIGS. **2-3B**. A function of the financial exchange is to generate data regarding events, such as financial market data events, e.g. price changes, changes in available quantity, availability of new orders, and transmit that data to the market participants. The users of the financial exchange may be tracking the occurrence of such events, and basing their buy/sell orders upon the occurrence thereof. The event may be a market data event (such as a price movement for an equity that a user is following), an event that is unique to the participant (such as a price movement for an equity the participant already owns, or an indication that prior offer to buy or sell has been transacted), completion of a calculation conducted by the participant (such as a calculation of a relationship between equities, or a calculation of a relationship between markets), or any other type of event that may be of interest to the participant, such as an event external to the market itself.

[0084] In some embodiments, a system for detecting user behavior may detect that messages or particular portions of one or more messages are received after other messages and/or particular portions of the one or more messages. Alternatively, the system may detect this user behavior based off a pattern by which the messages or portions of the messages are sent, the data that is included in the messages or portions of the messages in relation to the order in those messages or portions of messages are sent, or some combination thereof.

[0085] The data messages sent by a user may consist of two conceptual parts, as was described above: an event dependent portion and an event independent portion. The two portions may be delimited in relation to the occurrence, i.e. timing thereof, of an event (or the timing of when the user may have learned of the occurrence of the event), such as an event that may be of interest to a user of the system. The event may be receipt of information, receipt of the result of computations, or some other information of interest to the user. For example, a user of a financial market may seek to capitalize on information relevant to trading. The user may start to send the necessary information to complete a buy/sell order on the financial market. The user may intend to optimize the transport protocol, of the transport layer, by altering the order in which information is sent as part of their buy/sell order. For example, by optimizing the TCP protocol of the transport layer and/or the interaction of the application layer therewith, the user may break up their buy/sell order into packets, and then cause delivery of all, or most, of the packets that contain event independent information before generating and/or sending the packet containing the event dependent information. By controlling the delivery and order of packets, the user may send the majority of the packets in advance of receiving information that is relevant to the buy/sell order, thereby ensuring that when the event dependent information is sent the transmission of the event dependent packet is quicker than it would be otherwise.

[0086] The disclosed system may monitor network activity of the users of the system. Monitoring user network activity may comprise monitoring the types of messages received from the user, the number of messages sent by the user, the frequency or rate at which messages sent by the user are sent, or the relationship between messages sent by the user. A user's network activity may include the number and types of messages that the user sends. For example, a

user may begin the process of sending an Ethernet frame comprising a speculative trade order to attempt to hold a place in line, i.e. ensure that a receiving device routes the frame towards the destination ahead of other received frames which commenced sending slightly later, where the bulk of the frame is transmitted prior to the event and the FCS being altered post-event, or where the sender will have sufficient information to decide whether to 1) finish generating a valid Ethernet frame containing all or portions of one or more order entry messages, or 2) corrupt or otherwise cause the Ethernet frame or its contents which optimistically began sending to be dropped or discarded somewhere downstream. The user may be speculating on the outcome of an event and sending the Ethernet frame in anticipation of an opportunity, or in anticipation of the event occurring. If the event occurs as anticipated, the user then sends a correct FCS with the frame comprising their order having already been received, giving the user an advantage over other users who may have waited until the event occurred before sending in their orders. However, the user may then send an invalid FCS on the outgoing frame to invalidate the Ethernet frame if the opportunity does not materialize or the event does not occur, or if the event that occurs is not what the user anticipated. A system for detecting this type of behavior may comprise detecting an invalid Ethernet frame or detecting that the user sends multiple invalid Ethernet frames, e.g. beyond a particular threshold.

[0087] As discussed above, a user may be motivated to optimize, e.g. manipulate, the protocols in order to get their trade order to an exchange before others based on as much information as possible. Even a user with a fast connection may want to optimize the protocols to be sure they beat other users with similarly fast connections and reaction times, etc.

[0088] In some implementations of a system designed to detect the user behavior described above, there is the chance that some unintentional behavior, which may appear to be optimistically generated or otherwise as optimizing traffic as described, may be flagged while trying to detect actual optimized transmission or otherwise optimistic user messaging behavior. For example, TCP is designed assuming that there is a likelihood that packets will arrive out of order, so that if some packets of data are sent out of order they may still be recombined by the receiving node and ordered correctly. As such, there may be instances where a message is received out of order in such a way so that what appears to be event dependent data is only received after all of what appears to be event independent data. Similarly, the FCS is included in an Ethernet frame due to the possibility of the frame being corrupted in transit, e.g. transmission of one or more Ethernet frames may be initiated which are eventually terminated by invalid FCS's caused by problems in network transmission, or other issues other than a user attempting to optimize, e.g. manipulate, the Ethernet protocol. As will be described, some embodiments of the system may implement additional functionality to distinguish inadvertent from intentionally optimistic or otherwise optimized user behavior or otherwise mitigate the effects of the disclosed embodiments to avoid unduly penalizing a user for problems beyond their control. For example, message transmission behavior may be monitored for a period of time, such as one hour, one day, etc., e.g. data transmissions, number of packet per message, number of messages per packet, inter-packet transmission latency, retries, occurrence of errors, etc. may be recorded, for all or a subset of users in order to establish

a baseline, e.g. a static or dynamic statistical, or otherwise, model, for what may be considered normal or acceptable behavior, e.g. a historical average. This baseline may then be utilized to compare future data transmissions, inter-packet gap latencies, retries, errors, etc., to distinguish legitimate network communications problems from intentional optimistic or otherwise optimized behavior or otherwise define thresholds for number of errors, retries, resends, below which are considered legitimate and above which are considered intentionally optimized or subject to further review. For example, if during the monitoring period, it is observed for a given source that the customer reaction time between when a customer begins receiving a packet from the electronic trading system to the time when the customer begins sending out a packet to the electronic trading system in possible response to the receipt of the message from the electronic trading system (possibly referred to as the "tick to trade" latency), was previously 80 nanoseconds but later was observed to be 4 nanoseconds, and if the specific network interface has an absolute minimum serialization time of approximately 50-70 nanoseconds to transmit the minimum possible sized Ethernet frame, one may know that the customer must have started sending out their order message prior to having fully received the contents of the exchange message. Similarly, one could also monitor the network serialization latency measuring the time required for the customer's traffic to enter the exchange's network to the time that same packet finishes traversing through the exchange's network to reach its final destination, such as the window/lag/latency between packets, e.g. packets carrying portions of the same message, averages no less than 100 microseconds with a maximum value of 500 microseconds, that can serve as a baseline to later identify averages inter-packet latencies far below 100 microseconds, or peak latency far less than 500 microseconds, as indicative of optimistic or otherwise optimized behavior or otherwise meriting further review.

[0089] Optimistic user behavior may include optimizing the protocol at, for example, the data link layer. More particularly, optimizing the Ethernet protocol may include: 1) a user of the system continuously streaming Ethernet frames across one or more physical connections at or near the maximum transmission rate, regardless of whether it has begun to receive any messages from the exchange, and then, e.g. in a majority of cases, corrupting the sent Ethernet frames by sending an invalid FCS for the Ethernet frame or otherwise generating the traffic to ensure the packet is dropped, or 2) a user begins to send out a potentially to-be-corrupted Ethernet frame on the receipt of every incoming packet (before its full payload has been received). In the second example the user may have software and/or hardware coupled at the data link layer which determines, after having already started transmitting a packet, whether to correctly send the bits which make up the FCS of the Ethernet frame until the user receives additional information, which causes the user to determine to either correctly complete or otherwise corrupt or alter the outgoing Ethernet frame.

[0090] In an alternative embodiment, the system may look at the incoming TCP packets from a user and determine that, based on the TCP packets that have been received so far from the user, all of the independent portions have been received from that user. So, even though an event dependent portion has yet to be received, the system knows, because it

already has all of the event independent portions, that the event dependent portion must be the last portion when it is received so the system doesn't have to wait for the event dependent portion to make a decision, i.e. take an action, as described below, with respect to the packets or against the user.

[0091] While the disclosed embodiments may be discussed with respect to the TCP/IP and/or Ethernet protocols, it will be appreciated that the disclosed embodiments may be used with any network or communications protocol implemented at any physical and/or logical layer, particularly those which implement packet management, optimization and/or manipulation protocols, such as error correction or reliability protocols. For example, the disclosed embodiments, as opposed to being implemented with "best effort" protocols, may be implemented with respect to packet prioritizing or network resource conservation/reservation mechanisms, such as quality of service (QOS) protocols, which act, for example, to prioritize the transmission of some data over other data, e.g. based on source, data type, network congestion level, time of transmission, etc. It will be appreciated that such protocols may be managed/optimized, as described above, to intentionally optimize the delivery of messages, such as trade order messages, as described.

[0092] FIG. 2 depicts an illustrative environment in which the disclosed message processing system 210 may be implemented as described. In FIG. 2 the message processing system 210, which in some embodiments may be coupled with or be part of a receiver 206 of the electronic trading system 208, receives messages and transmits messages via a network 204a between sources 102a . . . n and the receiver 206 of the electronic trading system 208. As described herein, the receiver 206 may comprise a network switch, such as the ingress switch of electronic trading system 208, i.e. the first network component to receive incoming messages from the network 204a, or another network switch within the network infrastructure 302 of the electronic trading system 208, such as the last network switch prior to the orderer 310, or the orderer 310 itself, as will be described below. The receiver 206 may further comprise a component of the network interconnection or protocol stack of a switch or the orderer 310 as described above, e.g. the receiver 206 may be the software and/or hardware components which implement the protocol stack of a network switch or the orderer 310, etc.

[0093] Herein, the phrase "coupled with" is defined to mean directly connected to or indirectly connected through one or more intermediate components. Such intermediate components may include both hardware and software based components. Further, to clarify the use in the pending claims and to hereby provide notice to the public, the phrases "at least one of <A>, , . . . and <N>" or "at least one of <A>, , <N>, or combinations thereof" are defined by the Applicant in the broadest sense, superseding any other implied definitions here before or hereinafter unless expressly asserted by the Applicant to the contrary, to mean one or more elements selected from the group comprising A, B, . . . and N, that is to say, any combination of one or more of the elements A, B, . . . or N including any one element alone or in combination with one or more of the other elements which may also include, in combination, additional elements not listed.

[0094] FIG. 2 depicts exemplary computer devices, or sources, 102a, 102a, and 102n which may comprise different exemplary methods or media by which a computer device may be coupled with, either directly or indirectly, the message processing system 210 or by which a user may communicate, e.g. send and receive, trade or other information therewith. It will be appreciated that the types of computer devices deployed by market participants and the methods and media by which they communicate with while interacting with the message processing system 210 are implementation dependent and may vary and that not all of the depicted/described computer devices and/or means/media of communication may be used and that other computer devices and/or means/media of communications, now available or later developed may be used. Each computer device, which may comprise a computer 800 described in more detail below with respect to FIG. 8, may include a central processor that controls the overall operation of the computer and a system bus that connects the central processor to one or more conventional components, such as a network card or modem. Each computer device may also include a variety of interface units and drives for reading and writing data or files and communicating with other computer devices and while interacting with the devices that make up the message processing system 210. Depending on the type of computer device, a user can interact with the computer with a keyboard, pointing device, microphone, pen device or other input device now available or later developed.

[0095] An exemplary computer device 102a is shown coupled with the electronic trading system 208, via a network 204a such as via a T1 line, a common local area network (LAN) or other wired and/or wireless medium for connecting computer devices, such as the network 820 shown in FIG. 8 and described below with respect thereto. The electronic trading system 208 may further include a network infrastructure 212, such as the physical network layer 302 shown in FIGS. 3A and 3B, to interconnect the network 204a with the receiver 206. This infrastructure 212 may be implemented by a network such as the network 820 shown in FIG. 8 and described below with respect thereto. The interconnection/interface between the network 204a and the network infrastructure 212 of the electronic trading system 208, which may be referred to as an "edge" 204b, demarcates the internal network of the electronic trading system 208 from the external public network 204a, as described below. The user of the exemplary computer device 102a, or the exemplary computer device 102a alone and/or autonomously, may then transmit a trade or other information to the receiver 206 where the trade or other information is handled by the message processing system 210.

[0096] As was described above, the sources 102a . . . n, i.e. the market participants, may include one or more market makers which may maintain a market by providing constant bid and offer prices for a derivative or security to the electronic trading system 208, such as via one of the exemplary computer devices depicted. One skilled in the art will appreciate that numerous additional computers and systems may form part of the electronic trading system 208. Such computers and systems may include clearing, regulatory and fee systems.

[0097] The sources 102a, 102b, and 102n, transmit messages to the receiver 206 over a network 204a. These messages are received at the receiver 206 after they cross a network edge 204b, e.g. the logical demarcation between the

network **204a** and a network infrastructure **212**. Specifically, the network edge **204b** may be a logical boundary of the exchange **208**, which may also be referred to as the demarcation point or edge **204b**, may be defined as the first, or last, point at which the receiver **206**, or other network components controlled by the electronic trading system **208**, can control or otherwise optimize an incoming, or outgoing, transmission, e.g. data message or packet. For example, for an outgoing data packet, the edge **204b** of the exchange **208** may be defined as the last point at which the receiver **206**, or component thereof, can recall or otherwise stop the transmission. For example, the demarcation point or edge **204b** may be the point at which a market data message is provided to the protocol, e.g. multicasting or unicasting, for transmission or other point where data packets are individually forwarded toward their respective destinations, e.g. individually distinguishable by destination address. In at least one disclosed embodiment, the edge or demarcation point **204b** may further be defined as the point at which data messages or packets destined for multiple market participants are transmitted, e.g. multi-casted or otherwise, simultaneously, or substantially simultaneously, i.e. transmitted within short a time period such that an observer would consider them simultaneously transmitted or otherwise find the difference there between practically, logically or physically imperceptible. Thereafter, variation in network path latencies, etc. may impart unequal delays on the delivery of those messages.

[0098] Generally, the edge **204b** will lie between a component in the electronic trading system **208**, such as the receiver **206**, and another component controlled by another entity, such as an Internet Service Provider or other operator of another network. As described above, the edge or demarcation point **204b** may be geographically located anywhere, e.g. it may be geographically proximate to or remote from the receiver **206**. In some embodiments, market participants may collocate devices for receiving data from the receiver **206** in the same geographic location as the components of the receiver **206** which transmit that data.

[0099] The operations of computer devices and systems shown in FIG. 2 may be controlled by computer-executable instructions stored on a non-transitory computer-readable medium. For example, the exemplary computer device **102a** may include computer-executable instructions for receiving order information from a user and transmitting that order information to the receiver **206**. In another example, the exemplary computer device **102a** may include computer-executable instructions for receiving market data from the receiver **206** and displaying that information to a user.

[0100] Typically, messages from the sources **102a . . . n** to the electronic trading system **208** are handled in the order, or sequence, in which they are received. The message processing system **210** of FIG. 2, utilizing the disclosed embodiments, allows for moving away from handling messages in pure sequence (e.g. the Equilibrium model) and may reduce the value of being the very fastest user of a system with the smallest, e.g. absolute smallest, packet size. Moving toward the Equilibrium model may result in a more equitable system of handling messages and allow customers of the electronic trading system to focus on optimizing trading logic (offering or bidding better prices) rather than focusing resources on low level details of particular underlying messaging protocols or particular vendors' implementations

of those protocols which may trivially impact serialization time by a few nanoseconds but therefore lead to reduced queue spot.

[0101] Of course, numerous additional servers, computers, handheld devices, personal digital assistants, telephones and other devices may also be connected to the message processing system **210**. Moreover, one skilled in the art will appreciate that the topology shown in FIG. 2 is merely an example and that the components shown in FIG. 2 may include other components not shown and be connected by numerous alternative topologies.

[0102] In particular, example architectures, depicted in FIGS. 3A and 3B, for an electronic trading system **208** are disclosed which may be used with the disclosed embodiments. These architectures for the electronic trading system may be dependent on reliable handling of messages received by the electronic trading system.

[0103] In one implementation, as shown in FIG. 3A, the electronic trading system **208** may include an order processing function **314** which includes a receiver **206** which may comprise one or more network devices, such as switches, gateways, etc. which are coupled with a physical network layer **302** which interconnects the receiving device **206**, and generally the electronic trading system **208**, with the public network **204a** and market participant systems **102a-n**, etc. The receiver **206**, which may comprise an ingress switch to the electronic trading system **208**. Alternatively, the receiver **206** may comprise an orderer **310** as described herein, receives incoming transaction requests, e.g. incoming trade orders, and forwards them to a match engine **308**, e.g. via the physical network layer **302**, for matching as described herein. It will be appreciated that, for the purpose of this disclosure, the receiver may further include the last network component, e.g. switch, gateway, router, etc., which precedes the orderer **310**. Other components, not shown, may, based on the result of the matching process, formulate reporting or market data messages and transmit those messages to the market participants **102a-n**.

[0104] As will be described in more detail below with respect to FIG. 3B, an electronic trading system may further implement a set **306** of redundant match engines **308** to improve fault tolerance. This set **306** of redundant match engines **308** may include two or more match engines **308**, such as three or five match engines **308**. Incoming transactions, e.g. orders to trade, are received by the receiver **206** of the order processing function **314** which may further comprise an Orderer **310** component of the architecture which serializes, or otherwise sequences, the incoming transactions based on their order of receipt by the Orderer **310**. It will be appreciated that, for the purpose of this disclosure, the receiver may further include the last network component, e.g. switch, gateway, router, etc., which precedes the orderer **310**.

[0105] While the Orderer **310** of FIGS. 3A and 3B may be thought of as the point of determinism for the exchange **208**, as each transaction is then augmented with an indicium, such as a time stamp or other sequence encoding, indicative of its order of receipt relative to the other received transactions, ensuring their ordered processing thereafter, the receiver itself may actually be considered the point of determinism. That is so because as trade order messages, i.e. the frames/packets which comprise those messages, are received by the exchange **208** and routed/forwarded to the receiver **206** by physical network layer **302**, they ultimately converge at the

input of the receiver **206** where they are serialized in the order of receipt (or based on the arbitration protocol of the network device as was described above) and provided to the orderer **301** which then ascribes the time stamp to the essentially already-ordered transactions.

[0106] The sequenced transactions are then provided to the match engine **308**, as shown in FIG. 3A or substantially simultaneously communicated, e.g. broadcasted, to each match engine **308** of the set of redundant match engines **306**, in the implementation of FIG. 3B, each of which then processes the transaction, based on the sequencing imparted by the orderer, and determines a result, referred to as a match event, indicative, for example, of whether the order to trade was matched with a prior order, or reflective of some other change in a state of an electronic marketplace, etc. As used herein, match events generally refer to information, messages, alerts, signals or other indicators, which may be electronically or otherwise transmitted or communicated, indicative of a status of, or updates/changes to, a market/order book, i.e. one or more databases/data structures which store and/or maintain data indicative of a market for, e.g. current offers to buy and sell, a financial product, described in more detail below, or the match engines associated therewith, and may include messages which are indicative of, or otherwise generated based upon:

- [0107]** REST—indicates that a new order has been placed on an order book but not matched with a previously received order counter thereto (this event may also be indicated by a series of price improvement match events or deep book change match events, which may both be considered rests);
- [0108]** FILL—indicates that a new incoming order matched with one or more previously received but unsatisfied orders which were resting on an order book resulting in a trade;
- [0109]** MOD—indicates that an existing/resting order's values (price, quantity, etc.) have been modified/changed;
- [0110]** CANCEL—indicates that an existing/resting order has been canceled/removed;
- [0111]** MARKET OPEN—indicates that a market for trading has opened;
- [0112]** MARKET CLOSE—indicates that a market for trading has closed;
- [0113]** MARKET HALT—indicates that a market for trading has been paused for some period of time due to internal restrictions (usually that price velocity has gotten too high);
- [0114]** NEW PRODUCT—indicates that a new product is available;
- [0115]** CLOSE/CANCEL PRODUCT—indicates that a product is removed from trading;
- [0116]** PRODUCT TRANSITION—indicates that the market for a product is transitioning state, e.g. opening, closing, pausing, or reserving;
- [0117]** TRADING SCHEDULES—indicates the market hours;
- [0118]** FIRST TRADE—indicates that a first trade for a product has been placed;
- [0119]** PRODUCT LIMITS—indicates the price limits for a product;
- [0120]** TRADE—indicates that a trade for a particular product has occurred;

[0121] BUST—indicates that a trade has been invalidated;

[0122] RFQ—indicates a request for quote, e.g. a request to send in orders for a particular product;

[0123] HEARTBEAT—indicates an administrative message of the electronic trading system used to ensure communications of market events are functioning properly;

or other event or status.

[0124] FIGS. 3A and 3B show a block diagram depicting, in more detail, an example of the receiver **206** shown in FIG. 2, according to one embodiment. In one embodiment the receiver **206** may be a switch. In another embodiment the receiver **206** may be a TCP buffer. The electronic trading system **208** includes an interconnecting infrastructure, such as a physical communication network **302**, which may include network devices such as gateways, switches, and interconnecting media there between, such as backplane interconnects, optical and electrical communications media or other wired or wireless interconnect. The interconnecting infrastructure generally couples the various components of the electronic trading system **208** together and with market participant devices **102a . . . n**.

[0125] The electronic trading system **208** includes a match engine function embodied in the match engine **308**. In FIG. 3B the match engine function may be implemented by one or more sets **306** of redundant transaction processors **308**. While a single set **306** of match engines **308** will be described herein, as shown in FIG. 3B, it will be appreciated that many such sets **306** may be implemented both to improve fault tolerance through redundant operation and to increase the transaction handling capacity of the electronic trading system **208**, which encompasses an exchange.

[0126] Coupled with the match engine **308**, via the interconnecting infrastructure, is the order processing function **314** of the electronic trading system. In one embodiment, the order processing function is implemented on one or more FPGA devices, i.e. by one or more logic components thereof, coupled with the network gateway device (not shown), such as via a backplane interconnect, through which incoming transactions ingress the electronic trading system **208** and outgoing messages egress the electronic trading system **208**. The network gateway device, such as the receiver **206**, is further coupled with the interconnecting infrastructure to which the match engine **308** is also coupled. As shown in FIG. 3B, it will be appreciated that the set **306** of transaction processors may be coupled with the order processing function **314** via other means such as a dedicated interconnection there between. Further, as was discussed above, the disclosed mechanisms may be implemented at any logical and/or physical point(s) through which the relevant message traffic, and responses thereto, flows or is otherwise accessible, including one or more gateway devices, modems, the computers or terminals of one or more traders, etc.

[0127] The order processing function **314** receives incoming transactions from the market participants **102** and ensures deterministic processing thereof, i.e. that the incoming transactions are processed according to the defined business rules of the electronic trading system **208**, e.g. in the order in which they are received by the receiver **206**. Further, the order processing function (not shown) receives the output of the match engine **308** and evaluates that output to determine the correct result. The order processing func-

tion (not shown) may then further generate, or cause to be generated, appropriate acknowledgements and/or market data based thereon which are then communicated to the market participants 102.

[0128] FIG. 3B depicts an alternative illustrative electronic trading system that may be used to implement aspects of the disclosed embodiments. The system includes a physical network layer 302, a set of match engines 306, a plurality of match engines 308, an orderer 310, a decider 312, an exchange 208, and a receiver 206. The orderer 310 and/or decider 312 may be implemented as an FPGA and/or as part of the receiver 206.

[0129] More particularly, FIG. 3B depicts a block diagram of an exchange 208, which may also be referred to as an architecture, for processing a plurality, e.g. a series or sequence, of financial transactions, such as orders to trade a financial product, received via a network, from a plurality of market participants 102a . . . n, the processing of each transaction operative to cause a change in a current state of an electronic marketplace for one or more financial products. In one embodiment, each transaction may comprise a request to transact, e.g. an order to buy or sell, one or more financial products. A request to transact may further comprise a request to cancel a previous transaction, a status inquiry or other transaction.

[0130] The order processing function 314 of the exchange 208 includes a transaction receiver 310, e.g. an orderer as described above, which may be implemented as one or more logic components such as on an FPGA which may include a memory or reconfigurable component to store logic and processing component to execute the stored logic, such as via the interconnection infrastructure 302, and operative to receive each of the plurality of financial transactions and, upon receipt, augment, or otherwise ascribe or associate with, the received financial transaction with sequence data, such as an ordering or sequence, e.g. a sequence number or time stamp, indicative of a relationship, temporal or otherwise based on business rules/logic, e.g. a deterministic relationship, between the received financial transaction, e.g. the time of receipt thereof, and any of the plurality of financial transactions, e.g. the times of receipt thereof, previously and/or subsequently received by the transaction receiver 310. The ascribed ordering may then implicitly define the relationship with those transactions received thereafter. In one embodiment, the ordering may be an indicator of sequence, e.g. time stamp or, alternatively, an incremented sequence number.

[0131] The exchange 208 also includes a plurality 306 of transaction processors 308, e.g. match engines, coupled with the transaction receiver 310, such as via the communications infrastructure 302, each of the plurality 306 of transaction processors 308 operative to receive each of the augmented financial transactions and process, e.g. apply the business logic/matching algorithm to, the received augmented financial transaction in accordance with the sequence data to determine the change in the current state of the electronic marketplace caused thereby. As was described above, the processing is irrespective of the sequence in which each of the augmented financial transactions are received from the orderer, which may be different from the relationship indicated by the sequence data and which may result in a different change in the state of the electronic marketplace.

[0132] In one embodiment of the exchange 208, the processing of received augmented financial transactions imple-

ments a central limit order book of a financial market for at least one financial instrument.

[0133] In one embodiment of the exchange 208, each of the plurality 306 transaction processors 308 operates asynchronously with respect to the others of the plurality 306 of transaction processors 308, but, if operating properly, process the augmented financial transactions the, same, i.e. according to the sequence data and the applicable business rules. It will be appreciated that transaction processors 308 of redundant set 306 may be added or removed at will.

[0134] In one embodiment of the exchange 208 wherein the relationship indicated by the sequence data of a particular augmented financial transaction with respect to others of the augmented financial transactions is different from a relationship indicated by the order of receipt by one or more of the plurality of transaction processor of the particular augmented financial transaction with respect to the others of the augmented financial transactions, such as due to underlying processing priorities, transmission and/or routing anomalies, and would result in a different state change in the electronic marketplace.

[0135] In one embodiment of the exchange 208, each of the financial transactions comprises a request to transact in one of the one or more financial products, such as futures contracts or other financial instruments, e.g. swaps, equities, etc., the processing of each augmented financial transactions comprising identifying whether a previously processed augmented financial transaction remains incomplete and is counter thereto and, if so, indicating that a transaction there between may be completed, and if not, indicating that data indicative of the availability of the augmented financial transaction be stored in a database. It will be appreciated that while the disclosed embodiments are described with respect to the processing of financial transactions, the disclosed embodiments may be implemented with other transaction processing systems featuring or otherwise being characterized by a competitive deterministic behavior as described herein, such as systems which sell tickets or passes to events, e.g. concerts or sporting events, systems which sell unique or limited quantities of items, such as antiques, or other consumer goods/products having a restricted supply.

[0136] The order processing function 314 of the exchange 208 further includes a result arbiter 312, e.g. a decider as described above, which may be implemented as one or more logic components such as on the same or a different FPGA as the orderer 310, coupled with each of the plurality 306 of transaction processors 308, such as via the communications infrastructure 302, and operative to receive therefrom at least one of the determined changes in the state of the electronic marketplace for each processed augmented financial transaction and, based thereon, determine a selected change in the current state of the electronic marketplace for the processed augmented financial transaction and apply the selected change in the current state of the electronic marketplace to update the state of the electronic marketplace, the current state of the electronic marketplace now reflective thereof.

[0137] In one embodiment of the exchange 208, the transaction receiver 310 and result arbiter 312 are implemented in a network switch coupled with the data link layer/network layer of the communications infrastructure.

[0138] In one embodiment of the exchange 208, the result arbiter 312 is operative to compare the received determined changes in the state of the electronic marketplace for each

processed augmented financial transaction, and determine the selected change in the current state of the electronic market place to be the received determined change in the state of the electronic marketplace for each processed augmented financial transaction provided by, for example, the majority or a quorum of the plurality of transaction processors.

[0139] In one embodiment, the exchange 208, the result arbiter 312 may further determine that a transaction processor 308 of the plurality 306 of transaction processors 308 is faulty when the determined change in the state of the electronic marketplace for a processed augmented financial transaction received therefrom fails to agree with the he determined changes in the state of the electronic marketplace for a processed augmented financial transaction received from at least one other of the plurality 306 of transaction processors 308. The determination may be subject to a time delay threshold defining an amount of time which must elapse without having received a result before a fault is declared. As will be described, this threshold may be defined so as to prevent determination of a fault when a delayed result is expected, such as when a particular transaction processor 308 is known to be performing maintenance operations or is otherwise busy, offline or deactivated.

[0140] For example, in one embodiment of the exchange 208, each of the plurality 306 of transaction processors 308 is operative to periodically perform one or more other functions, such as maintenance, e.g. garbage collection, during which augmented financial transactions are not processed or processing is delayed. In this embodiment, each of the plurality 306 of transaction processor 308 may be further configured to not perform the one or more other functions contemporaneously with the performance of the one or more other functions by the remaining of the plurality 306 of transaction processors 308. Alternatively, more than one transaction processor 308 may be allowed to perform other operations assuming a sufficient number are remaining to meet a requisite level of fault tolerance.

[0141] In one embodiment of the exchange 208, the plurality of financial transactions may further include a plurality of administrative transactions, each of which may or may not cause a change in the current state of the electronic marketplace. Such administrative transactions may include instructions to configure the transaction processors 308, such as to synchronize their operation or cause them to perform maintenance or other operations.

[0142] As described above, the financial transactions processed by the system depicted in FIGS. 3A and 3B may all be received as messages, or data messages. The financial transactions, and their corresponding parts, or portions, may all be sent through the network, including the physical network layer 302, and the messages will at one point be handled by the receiver 206, or gateway.

[0143] FIG. 4A depicts an illustrative implementation of the message processing system 210, according to the disclosed embodiments for processing one or more data messages received via a network from a source 102a, each, or a subset thereof, data message comprising a portion that is generated by the source dependent upon an event and a portion that is generated by the source independent of the event. As shown in FIG. 4A, the message processing system 210, which may be part of a receiver 206, includes a network interface 402, a processor 404, memory 406, instructions 408, and a buffer 410. The message processing system 210 may receive data

messages from the sources 102a, 102b, and 102n over a network 204a and network infrastructure 212. The data messages may be transmitted in portions that may be dependent on an event and independent of an event. The data messages may be received at the network interface and may be processed by the processor 404 in conjunction with the memory 406.

[0144] FIG. 4B depicts a functional implementation of the system 210 of FIG. 4A. As shown in FIG. 4B, the message processing system 210, which may be part of a receiver 206, includes a network interface 402 connected to a network 204a via a network infrastructure 212, a message receiver 412, a message analyzer 414, a message processor 416, and a buffer 414. The network interface 402, which may comprise the network interface 818 described below with respect to FIG. 8, is operative to couple the message system 210 with a network, such as a public and/or private communications wired and/or wireless network, such as the network 820 shown in FIG. 8 and described in more detail below, and facilitates communications among the message processing system 210 and the participants 102a, 102b, and 102n as described herein.

[0145] The message receiver 412, which may be implemented by the processor 404 and memory 406 shown in FIG. 4A, or as a separate component or as one or more logic components, such as on an FPGA which may include a memory or reconfigurable component to store logic and a processing component to execute the stored logic, or as first logic, e.g. computer program logic, 502, described below, stored in the memory 406, which may be implemented by a memory 812 shown in FIG. 8 and described in more detail below with respect thereto, or other non-transitory computer readable medium, and executable by a processor 404, which may be implemented as a processor 804 shown in FIG. 8 and described in more detail below with respect thereto, to cause the processor 404 to, or otherwise be operative to receive a data message, as described above. The message receiver 412 may be further operative to receive the data message from a network switch. For example, the message receiver 412 may be a TCP buffer of the computer which is the destination of the message, such as the orderer 310 described above.

[0146] The message analyzer 414, which may be implemented by the processor 404 and memory 406 shown in FIG. 4A, or as a separate component or as one or more logic components, such as on an FPGA which may include a memory or reconfigurable component to store logic and processing component to execute the stored logic, or as second logic, e.g. computer program logic, 504, described below, stored in the memory 406, which may be implemented as a memory 812 shown in FIG. 8 and described in more detail below with respect thereto, or other non-transitory computer readable medium, and executable by the processor 404, which may be implemented as a processor 804 shown in FIG. 8 and described in more detail below with respect thereto, to cause the processor 404 to, or otherwise be operative to, automatically analyze the data message to determine if the event independent portion was transmitted by the source prior to occurrence of the event by evaluating whether all of the event independent portion was received prior to any of the event dependent portion. The message analyzer 414 may automatically analyze the data upon receipt of the data message, e.g. in real time, or it may analyze the data at a later time. In some embodiments, the message processing system 210 may evaluate the data

message to see if the data message could have been sent in a single packet but was instead sent by the source in multiple packets. This evaluation may be performed in conjunction with and prior to looking at the order of the portions received (the multiple packets). For example, a source could send some event independent data late along with the event dependent data.

[0147] The message analyzer 414 may be further operative to calculate a minimum size for a complete data message, and further calculate a product of the size and a serialization time for the data message through a particular physical layer. The calculated values may be used to penalize data messages or portions thereof, e.g. delay the data message by the calculated amount of time before providing the message to the application, e.g. the orderer. In other embodiments, the message analyzer 414 may be further operative to calculate a variance in serialization time for the data message. A variance in serialization time for a data message may be a deviation in the amount of time it takes to transmit a data message from an expected amount of time it takes to transmit a data message. In some embodiments, the message analyzer 414 determines when a source of a data message exceeds a network activity level threshold. Alternatively, or in addition thereto, the delay may be computed based on variance in transit time from the message source 102a to the message receiver 412. For example, a first message source may be located, logically and/or physically, remotely from the system 210 such that message transmission time there between averages 1 millisecond (ms), while a second message source may be located remotely from the system 210 such that the message transmission time averages 2 ms. Accordingly, an added delay of 0.5 ms would impact the first message source more so than the second message source. Therefore, a more substantial delay may be assigned to the second message source as a calculated delay, where appropriate, so as to have a proportional impact thereon.

[0148] In one embodiment, a normalized delay factor may be computed based on the latency of the slowest received data messages over, for example, a particular period of time, e.g. a trading period, whereby an apportioned amount the normalized delay may be applied to each data message, or completing portion thereof, to normalize the timing thereof with respect to all of the received data messages.

[0149] More particularly, in a situation where the data message is a TCP data message, or TCP packet, the message analyzer 414 may be further operative to determine which portions of the TCP packet have been sent by the source 102a and in what relation to the other portions. For example, in some instances the source 102a may send the first four packets of a TCP message but withhold the fifth packet of the TCP message dependent on an event that the source 102a is waiting to receive information about. In another instance the source 102a may send packets two through five thousand and withhold packet one until the source 102a receives information about an event before sending packet one. Both scenarios would involve an optimization relying on the reliability and/or ordering aspects of the TCP protocol, and potentially a violation of business rules of an Exchange if the source 102a is a participant in an electronic trading system. Both of these scenarios may involve a source 102 taking advantage of a situation where the software program and/or hardware implementing TCP on the destination computer has been programmed to request only a missing packet in a packet loss situation or where the application layer program

is awaiting a validation/checksum value to validate the portion of a message already received. For example, in the second scenario, the receiving device may receive packets two through five thousand and then the receiving device may respond as if packet one was lost, thereby requesting packet one be resent from the source 102a. In reality, packet one was not lost and was being withheld by the source 102a. In this situation, the disclosed embodiments, which determine the timing of the sending of event dependent and event independent portions of a data message, detect this type of optimization. As an alternative to counteract this optimization of packet loss handling, the TCP software/hardware may be programmed to request that a sender resend all of the packets that make up a data message. Additionally, the TCP software/hardware on the destination computer may be programmed so that when such optimization of TCP is detected the source 102a of the optimization may be identified and then how lost packets from that source 102a are handled by the software and/or hardware may be decided; resend only the perceived lost packet, or send all packets when a lost packet is perceived. Additionally, source 102a's TCP implementation could even go so far as to eliminate the retransmit and out of order segment capabilities of TCP and simply tear down any customer TCP connections which produce segments that are not transmitted in order or demonstrate dropped packets either on egress or ingress.

[0150] Similarly, the message analyzer 414 may be operative to monitor the received packets of the TCP message to determine if the source 102a is sending a portion of information that is dependent on an event prior to the rest of the event dependent portions being transmitted after all the event independent portions of the TCP message are transmitted. Put another way, a source 102a may send nine of ten event independent portions of a data message, then send one of three event dependent portions, then send the final event independent portion of the data message, and finish the transmission with the two remaining event dependent portions which are only sent after the event occurs.

[0151] In another example, where the data message is an Ethernet frame, the source 102a may start to submit an Ethernet frame, the payload of which comprises, for example, an event dependent portion, prior to the occurrence of an event/trigger. For example, a source 102a may start submitting an Ethernet frame, i.e. comprising a trade order, as the source 102a receives information indicative of the occurrence of some not yet fully known/appreciated event, i.e. that might constitute a trigger. When that information does not constitute a trigger, the participant invalidates the outbound Ethernet frame (by e.g. sending an incorrect FCS). This behavior may be detected by the message analyzer 414 by determining if a source 102a is continuously streaming in Ethernet at or near a maximum rate and corrupting the packets, or if a source 102a begins to send out a potentially to-be-corrupted packet on the receipt of every incoming packet i.e. before its full payload has been received. The message processing system 210 may monitor a source's network activity level to determine streaming rates of Ethernet frames, as well as other Ethernet frame activity.

[0152] In further embodiments, market data feeds may be provided by the Exchange 208. The market data feeds may transmit market data messages to sources 102a . . . n indicative of market events and which may constitute and event/trigger upon which the recipients may wish to transact. Accordingly, the message processing system 210 may

monitor outgoing market data feeds (or other events within the exchange's control) and determine that a data message has been received when a portion of data message related to an outgoing market data feed message (related via the same financial instrument for example), is received prior to the transmission of the associated market data message, and the remainder is received after transmission. Alternatively, the message processing system 210 may determine that the remainder is received after transmission of the market data message but before the market data message could have reasonably been received and/or processed by the source 102a, such as based on known/measured or estimated latencies between the message processing system 210 and the source 102a. The determination may be done by the message analyzer 414, or another component, such as an FPGA, or a network switch.

[0153] The message processor 416, which may be implemented by the processor 404 and memory 406 shown in FIG. 4A, or as a separate component or as one or more logic components, such as on an FPGA which may include a memory or reconfigurable component to store logic and processing component to execute the stored logic, or as third logic, e.g. computer program logic, 506 stored in the memory 406, which may be implemented as a memory 812 shown in FIG. 8 and described in more detail below with respect thereto, or other non-transitory computer readable medium, and executable by the processor 404, which may be implemented as a processor 804 shown in FIG. 8 and described in more detail below with respect thereto, to cause the processor 404 to, or otherwise be operative to, when it is determined that the event independent portion was transmitted by the source 102a prior to occurrence of the event, take an action with respect to the data message. For example. The message processor 416 may be operative to delay the forwarding of the message to the Exchange, e.g. to the orderer 310 or otherwise cause the orderer to ascribe a later timestamp to the message so as to penalize the message as compared to other received messages. Such delays may be implemented via a storage buffer with a timed release, via a long path such as an appropriate length network path (e.g. physically longer cable), or other known or later developed methods, or combinations thereof, of imparting latency to the transmission of a message via a network. The message processor 416 may be further operative to disconnect a connection to the network for the source 102a of the data message when the data message is one of many received data messages received from the source and a network activity level for the source 102a, as determined by the many received data messages received from the source 102a, exceeds a network activity level threshold as determined by the message analyzer 414. For example, the message processing system 210 may keep track of a number of received messages of a particular protocol type (e.g. TCP, Ethernet, FIX) from a source over a period of time and determine that the source is abusing the message processing system 210. In some embodiments, the message processor 416 is further operative to add the variance, as calculated by the message analyzer 414, to a time of receipt for the data message. The message processor 416 may transmit the data message to a network switch, a financial market, or a buffer.

[0154] FIG. 7 depicts a flow chart showing operation of the system of FIGS. 1-4B according to one embodiment. In particular FIG. 7 shows a computer and/or FPGA implemented method for processing a plurality of messages

received via a network, such as the network 820 of FIG. 8, from a plurality of sources 102a . . . n, the processing of each transaction operative to process one or more data messages received via a network from a source, each, or a subset thereof, data message comprising a portion that is generated by the source dependent upon an event and a portion that is generated by the source independent of the event. In one embodiment, each data message may comprise data indicative of a message type, at least one message characteristic, and a message payload. In another embodiment, each portion of the received data message includes a time of receipt affixed by the message receiver which is used by the message analyzer to determine the order of the received portions of the data message. The source may correspond to at least one of a plurality of market participants. Similarly, the message type may be a TCP/IP packet, a FIX message, or an Ethernet frame. In some embodiments, the data message received may be an incomplete data message and the at least one message characteristic of the data message includes a level of completeness for the data message. Similarly, in some embodiments, the incomplete data message has data indicative of fields that comprise the data message, and wherein the level of completeness includes data indicators indicating which fields of the incomplete data message contain data. The data fields of the data message may correspond to the fields used in the TCP, IP, Ethernet, or FIX protocols. The data indicators may indicate when data is in each field, and what data is located in each field.

[0155] The flowchart 700 includes receiving, by a message receiver 412 from a network, such as via the interconnection infrastructure 302 of FIG. 3A or 3B or the network interface 402 of FIGS. 4A and 4B, a data message (Block 702). The data message may be received from a network switch, from a source or market participant, or from a buffer of incomplete data messages, such as the buffer 410 described in FIG. 6 below.

[0156] The flowchart 700 further includes analyzing, by the message analyzer 414, coupled with the message receiver 412, such as via the communications infrastructure 302 of FIG. 3A or 3B or the network interface 402 of FIGS. 4A and 4B, the data message to determine if an event independent portion was transmitted by the source prior to occurrence of an event (Block 704). As was described above, analyzing may include analyzing the data message to determine if the event independent portion was transmitted by the source prior to occurrence of the event by evaluating whether all of the event independent portion was received prior to any of the event dependent portion. Similarly, the analyzing of the data message may include calculating a size for the data message, calculating a product of the size and a serialization time for the data message. In other embodiments, the analyzing may include calculating a variance in serialization time for the data message. In still more embodiments, the analyzing determines when a source of a data message exceeds a network activity level threshold.

[0157] The flowchart 700 further includes taking action, by a message processor 416, coupled with the message analyzer 412, such as via the communications infrastructure 302 of FIG. 3A or 3B or the network interface 402 of FIGS. 4A and 4B, with respect to the data message (Block 706). In some embodiments, where the received data message is an incomplete data message taking an action comprises taking an action for the incomplete data message based on the level of completeness for the data message. In other embodi-

ments, taking an action further comprises disconnecting a connection to the network for the source of the data message when the network activity level for the source exceeds a network activity level threshold as determined by the message analyzer. Similarly, taking an action with respect to a data message further comprises applying a delay to the data message as described above. The delay may include delaying forwarding the data message. Delaying forwarding may comprise delaying the forwarding of the data message based on the serialization of the data message, i.e. how long it takes the data message to be transmitted. The delay may be computed based on a size of the data message, calculating a product of the size and a serialization time for the data message, and the message processor being further operative to add the product to a time of receipt for the data message. Serialization time for the data message may be the amount of time it takes for the data message to be transmitted across a network based on the data message. For example, a delay dependent on the size of the data message, which would equal the time it would take to submit the data message on the lowest bandwidth infrastructure component between participant and gateway. For example, if participants connect via a 10 Gb connection and a buffer contains 200 bytes, the delay would be (200 bytes×approximately 1 nanosecond serialization time per byte on 10 Gb=approximately 200 nanoseconds). A meaningful delay may be dependent on the variance of the network transport time. For example, a delay of one millisecond could be added, which for co-located machines with usual latencies of <1-10 microseconds, should more than offset any benefit resulting in splitting order traffic.

[0158] Similarly, taking action with respect to a data message may further comprise transmitting the data message to a network switch, a financial market, or a buffer subsequent to taking the action.

[0159] FIG. 5 depicts an illustrative set of instructions for use with the systems of FIGS. 1-4B according to one embodiment. The instructions 408 may include computer code, such as the instructions 810 disclosed in FIG. 8 below, which implements the functionality of the systems disclosed in FIGS. 1-4B. In particular, the instructions 408, may include logic, or a first logic component, for a message receiver 502, further logic, or a second logic component, for a message analyzer 504, and yet further logic, or a third logic component, for a message processor 506. The instructions may be coupled with the network interface 402, and/or the buffer 410.

[0160] FIG. 6 depicts an illustrative buffer for use with the system of FIGS. 1-5 according to one embodiment. The buffer 410 may be a temporary storage area whose purpose is to act as a holding area for data that is received by the systems of FIGS. 1-5. The buffer 410 may hold data messages that are not yet ready to be processed, or transmitted. The buffer 410 may also hold portions of data messages that have been received. The portions of data messages held in the buffer 410 may include portions that are dependent and independent of events, or triggers, that are important to the sources, or market participants. The buffer 410 may be stored in a temporary memory such as Random Access Memory, or any of the memory shown in FIG. 8, such as the memory 812, or computer readable medium 808. More particularly, the buffer 410 may hold portions of TCP data messages, or portions of Ethernet frames, or both. The buffer

410 may be a source of data messages and a destination for data messages or their portions.

[0161] One skilled in the art will appreciate that one or more functions/modules described herein may be implemented using, among other things, a logic component such as a reconfigurable logic component, e.g. an FPGA, which may include a logical processing portion coupled with a memory portion, or as a tangible computer-readable medium comprising computer-executable instructions (e.g., executable software code) executable by a processor coupled therewith to implement the function(s). Alternatively, functions/modules may be implemented as software code, firmware code, hardware, and/or a combination of the aforementioned. For example the functions/modules may be embodied as part of an electronic trading system 208, which may include a message processing system 210, for handling messages related to financial instruments.

[0162] The following are exemplary scenarios to which the above described embodiments may be applicable:

Serialization Reduction/Packet Optimization Examples

[0163] A. Methods of commencing transmission but later getting packets discarded (at switch, router, NIC, TCP stack, application layer)

[0164] 1. Generate an invalid frame or message utilizing Ethernet FCS, IP checksum, TCP checksum, or FIX checksum

[0165] 2. Bring down or even optically block the physical layer connection (turn off transceiver) before full packet is transmitted

[0166] 3. Corrupt at 64/66 encoding level

[0167] 4. Send a frame with a larger than supported MTU, causing the switch or router to discard the message, possibly multiple hops downstream, before it reaches the gateway

[0168] 5. Send a (checksum) valid frame which is marked as repeat of previously sent TCP traffic on same session (so it will be discarded by the TCP stack invisibly to the application) or with prior sequence number (so it will be rejected by application).

[0169] 6. Send FIX message with either invalid tag, valid tag with invalid value (nonexistent contract), or other parameter that would pass through the FIX "session" level validation (valid checksum, contains valid next sequence number, msg type, etc) but fail at the business validation level for more complex reasons. See below table for list of example reject codes.

TABLE

Example Reject Codes	
Text explanation of error code	
Sent in response to New Order (tag 35-MsgType = D) and Cancel Replace (tag 35-MsgType = G) messages in tag 103-OrdRejReason of the Execution Report (tag 35-MsgType = 8)	
1003	Orders may not be entered while the market is closed
1003	Orders may not be entered while the market is paused
1012	Price must be greater than zero
1013	Invalid order qualifier

TABLE-continued

Example Reject Codes	
Text explanation of error code	
1014	The user is not authorized to trade
2013	Market price orders not supported by opposite limit
2019	Order's GTD Expire Date is before the current (or next, if not currently in a session) trading session end date
7000	Order rejected
2045	This order is not in the book
2046	Disclosed Quantity cannot be greater than total or remaining qty
2047	Order contract is unknown
2051	The Order was submitted with a different side than the requesting Cancel
2058	Stop price maxi-mini must be greater than or equal to trigger price
2059	Stop price maxi-mini must be smaller than or equal to trigger price
2060	Sell order stop price must be below last trade price
2061	Buy order stop price must be above last trade price
2100	The modify was submitted on a different product than the original order
2101	Attempt to modify an order with a different in-flight-fill mitigation status than first modification
2102	Attempt to modify an order with a different SenderCompID than the original order
2115	Order quantity is outside of the allowable range
2130	Order type not permitted while the market is in Post Close/Pre-Open (PCP)
2137	Order price is outside the limits
2179	Order price is outside bands
2311	Order type not permitted for group
2500	Instrument has a request for cross in progress
2501	Order Quantity too low
7024	Order cannot be modified or cancelled while the market is in no cancel
7027	Order type not permitted while the market is reserved
7028	Order session date is in the past
7029	Orders may not be entered while the market is forbidden
7613	Disclosed quantity cannot be smaller than the minimum quantity
Sent in response to Order Cancel Request (tag 35-MsgType = F) message in tag 102-CxlRejReason of the Order Cancel Reject (tag 35-MsgType = 9)	
1003	Orders may not be canceled while the market is closed
2045	This order is not on the book.
2048	The order was submitted with a different SenderCompID than the requesting cancel
7024	Order cannot be modified or cancelled while the market is in No Cancel

TABLE-continued

Example Reject Codes	
Text explanation of error code	
Sent in response to the New Order Cross (tag 35-MsgType = s) message in tag 103-OrdRejReason of the Execution Report (tag 35-MsgType = 8)	
2502	CrossType ineligible in Instrument
2503	RFQ for Cross is not allowed since contract is not RFQ for Cross eligible
3000	RFQ rejected, instrument/exchange not tradable
3001	RFQ rejected, instrument/exchange no known
7610	RFC not allowed during RTH.
Sent in response to Security Definition Request (tag 35-MsgType = c) message	
3002	Error creating contract
3010	Business Reject: Other
3011	Business Reject: Unknown ID
3012	Business Reject: Unknown SECURITY
3013	Business Reject: Unsupported Message Type
3014	Business Reject: Application Not Available
3015	Business Reject: Required Field Missing
3016	: Contract claims to be a spread but its missing one or more legs
4000	Engine did not respond to request
7005	Order contract is unknown
7009	The contract for this order is past expiration date and may no longer be traded
7009	The contract for this order is has a future activation date and cannot yet be traded
7010	Invalid spread price
7011	The maximum amount of working orders for this contract has been exceeded
7012	Submission of this order would cause an overwrite of an existing order on book
7013	Order group does not match group of contract
7014	Order Security type does not match security type of contract
7015	Order modify has different side than existing order
7020	No Trading Calendar found on or after Order's GTD Expire Date
7021	Tag ExpireDate (432) beyond instrument expiration
7022	Tag ExpireDate (432) beyond UDS COMBO earliest leg expiration
Sent in Response to Mass Quote (tag 35-MsgType = i) message in tag 300-QuoteRejectReason of the Quote Acknowledgment (tag 35-MsgType = b)	
2600	MarketMaker protection has tripped
7101	Quote rejected: Unknown instrument
7102	Quote rejected: Exchange (Security) closed
7103	Quote rejected: Quote exceeds limit
7104	Quote rejected: Too late to enter
7105	Quote rejected: Unknown Quote
7106	Quote rejected: Duplicate Quote in same message

TABLE-continued

Example Reject Codes	
	Text explanation of error code
7107	Quote rejected: Invalid bid/ask spread
7108	Quote rejected: Invalid price
7109	Quote rejected: Not authorized to quote this instrument
7110	Quote rejected, resting cancelled: Too many rejects
7111	Quote rejected: Duplicate Quote, different ID
7112	Quote rejected, resting cancelled: Quote exceeds limit
7113	Quote rejected, resting cancelled: Duplicate Quote
7114	Quote rejected, resting cancelled: Invalid bid/ask spread
7115	Quote rejected, resting cancelled: Invalid price
7116	Quote rejected, resting cancelled: Not authorized to quote this instrument
7117	Quote rejected: Unspecified reason
7118	Quote rejected: Exchange (Security) is in pre open
7119	Quote rejected: Exchange (Security) is in no cancel
7120	Quote rejected: Exchange (Security) is reserved
9999	Technical error - Function not performed.

[0170] B. Methods of avoiding corruption and sending another message instead

[0171] 1. Modify the “nature” of the packet’s FIX business message to avoid having to corrupt (instead of sending on the contract actually of interest or greatest importance, instead send a message on a different contract or at a different price which has a lower probability of being filled or different contract which is allowed to transact at that time)

[0172] 2. Modify the type of packet (instead of sending TCP packet, modify to be ICMP) or TCP packet to a different (or nonexistent) host

[0173] 3. Send smaller TCP disconnect packet on different fiber which may get to gateway first, tearing down the connection, and possibly (depending on application implementation) causing the application not to process the subsequent in transit complete messages which may be of longer length and therefore reach the gateway later than the shorter sized TCP disconnect message.

[0174] C. Examples of optimization or optimistic messaging while still using valid order messages and packets

[0175] 1 Optimize and minimize the length of any or all configurable variable length fields in FIX message (for example ensuring Tag 50, which must be unique at clearing firm level, is only two bytes in length).

[0176] 2. Exclude any optional FIX fields (HandInst, CustOrderHandlingInst, Symbol, TimeInForce, MinQty, SelfMatchPreventionInstruction, GiveUpFirm, CmtaGiveupCD, CorrelationClOrdID)

[0177] 3. Minimize the number of messages sent on particular session (in order to keep sequence number only a single byte); If sequence numbers reset weekly and participating in preopen for example, may require multiple FIX sessions, one or more used per day potentially. To demonstrate, the very first message of a FIX session will have a serial number of “1” and therefore that FIX tag will only be one byte long, but ten thousandth message sent on that FIX

session (before required reset) would be “10000” and therefore be five bytes long, increasing the sending time.

[0178] 4. Ensuring minimum possible header size for all protocols with potentially variable length (Ethernet, IP, TCP) excluding any optional headers or extensions which would increase packet size and therefore serialization time.

[0179] 6. Sending packets at faster than minimum inter-packet gap rate (if in microbursts where dropping may not occur) for those instances where exchange’s particular switches correctly handle and support smaller than specification-mandated IPG

[0180] 7. Sending packets with smaller than specified preamble (four vs eight bytes)

[0181] 8. In order to send N messages, utilize N separate fiber optic connections to exchange, sending a single message down each in parallel.

[0182] 9. Utilize alternative but nearly equivalent methods of achieving same business intent but utilizing the method with the smallest message size/type: Examples include submitting new orders away from the BBO, and then modifying those existing orders to aggress vs the alternative intuitive method of simply submitting a new order (which may be a longer message and therefore slower to transmit over network than shorter modify message of an existing order). Another example (depending on how exchange implements) would be a potential preference for utilizing TCP session disconnect message combined with a Cancel on Disconnect feature vs explicitly sending Fix cancel messages to cancel individual orders or utilizing Mass Cancel Quote, all of which would have longer serialization times than a message simply tearing down the session at the TCP level.

[0183] 10. Running clock speed of Ethernet TX at fastest possible clock speed that exchange’s switch or Ethernet device will or may sometimes still be able to process (basically overclocking the 10 Gb link and relying on the fact that the receiving logic is recovering or attempting to recover the remote transmitter’s transmit clock before passing the data through a dual clock FIFO into another slower clock domain running closer to the actual specified clock).

[0184] Alternatively, or in addition to the functionality described above, the disclosed embodiments may:

[0185] Ensure there are not multiple paths (whether dynamic or statically assigned) for different customers to reach the same gateway or point of ordering, to mitigate systemic latency disadvantage to particular customers or sessions;

[0186] Ensure switches and routers utilize only deterministic non-variable lookup data structures (TCAM not binary trees) to avoid different switches having different latencies or even different source/destination IPs having different latencies (otherwise customers could take actions which cause these data structures to be updated and then decrease or increase latency for particular paths or switches);

[0187] Strip out or ensure networking fabric ignores all customer provided QoS flags at various protocol levels, for example Ethernet (802.1Q or 802.1P), IP (DSCP), etc;

[0188] Have per-switch-port “latency normalization” (taking into account and correcting for differences in cable lengths, PCB traces on the switches themselves, temperatures, curvature of cables, local oscillators, etc.);

[0189] Provide “Firm aware” switch and router algorithm when choosing among multiple packets received at exactly the same time (if desired that a single customer with N fiber optic cables which sends N messages all in parallel, one

message down each fiber, doesn't have an advantage over another customer with only a single fiber optic cable which can only send a single message at a time);

[0190] Provide Enhanced business logic for specifically handling case where multiple customers are submitting messages at exactly the same time (same clock): options include pro-rata, prioritizing based on which messages will maximize revenue for the exchange, prioritizing based on which contracts, order sizes, etc.

[0191] With electronic trading systems which may utilize GPS or other time synchronization receivers which utilize PTP or PPS or similar to sync the "absolute time" to one or more switches, orderers, gateways, the disclosed embodiments may: ensure when doing time synchronization that time cannot go backwards to prevent previous later packet from being pulled into application earlier than earlier packet if time sync occurred after first packet but before second packet which indicated that the clock was "too far ahead" and now needed to slow down or halt incrementing temporarily; Mitigate GPS spoofing to alter time and trigger the above (for example if time is far off but cannot go negative, what if multiple packets get same timestamp while time is prevented from going forward); need separate non-synced timestamp that is based off a monotonically increasing counter clocked off of the recovered network clock and always monotonically increasing; Provide two timestamps, one which is synced to "absolute time" and a second which is just a monotonically increasing clock specifically recovered from the RX of the Ethernet link. This way it preserves the interpacket gap between messages to detect queuing, know which packets arrived first, etc.

[0192] In one embodiment, the electronic trading system may: provide a clock output to actually drive customers transmitting logic (to prevent customers from attempting to overclock or arbitrage exact clock frequencies of individual network links or switches); may Ethernet-level padding of frames may be utilized (or simply holding on to them at FPGA for a calculated amount of time) to ensure that differently formatted frames (at the various protocol levels) doing the same business-level operation (sending a buy order) have consistent latency (regardless of usage of optional IP or TCP extensions or utilization or non-utilization of optional FIX headers, longer length of variable length fields, etc).

[0193] First network hop on the exchange's side which is directly connected to the customer could inject timestamp of SOF (start of frame) and then provide that upstream to orderer so that packets of trivially smaller size are not given an advantage if they started being sent at exactly the same time as packets of slightly larger size (especially true if all packets being sent upstream in parallel via cut through switching or similar).

[0194] Referring to FIG. 8, an illustrative embodiment of a general computer system 800 is shown. The computer system 800 can include a set of instructions that can be executed to cause the computer system 800 to perform any one or more of the methods or computer based functions disclosed herein. The computer system 800 may operate as a standalone device or may be connected, e.g., using a network, to other computer systems or peripheral devices. Any of the components of the electronic trading system 208 discussed above may be a computer system 800 or a component in the computer system 800. The computer system 800 may implement a match engine, margin pro-

cessing, payment or clearing function on behalf of an exchange, such as the Chicago Mercantile Exchange, of which the disclosed embodiments are a component thereof.

[0195] In a networked deployment, the computer system 800 may operate in the capacity of a server or as a client user computer in a client-server user network environment, as a peer computer system in a peer-to-peer (or distributed) network environment, or as a network device such as a switch, gateway or router. The computer system 800 can also be implemented as or incorporated into various devices, such as a personal computer (PC), a tablet PC, a set-top box (STB), a personal digital assistant (PDA), a mobile device, a palmtop computer, a laptop computer, a desktop computer, a communications device, a wireless telephone, a land-line telephone, a control system, a camera, a scanner, a facsimile machine, a printer, a pager, a personal trusted device, a web appliance, a network router, switch or bridge, or any other machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine. In a particular embodiment, the computer system 800 can be implemented using electronic devices that provide voice, video or data communication. Further, while a single computer system 800 is illustrated, the term "system" shall also be taken to include any collection of systems or sub-systems that individually or jointly execute a set, or multiple sets, of instructions to perform one or more computer functions.

[0196] As illustrated in FIG. 8, the computer system 800 may include a processor 802, e.g., a central processing unit (CPU), a graphics processing unit (GPU), or both. The processor 802 may be a component in a variety of systems. For example, the processor 802 may be part of a standard personal computer or a workstation. The processor 802 may be one or more general processors, digital signal processors, application specific integrated circuits, field programmable gate arrays, servers, networks, digital circuits, analog circuits, combinations thereof, or other now known or later developed devices for analyzing and processing data. The processor 802 may implement a software program, such as code generated manually (i.e., programmed).

[0197] The computer system 800 may include a memory 804 that can communicate via a bus 808. The memory 804 may be a main memory, a static memory, or a dynamic memory. The memory 804 may include, but is not limited to computer readable storage media such as various types of volatile and non-volatile storage media, including but not limited to random access memory, read-only memory, programmable read-only memory, electrically programmable read-only memory, electrically erasable read-only memory, flash memory, magnetic tape or disk, optical media and the like. In one embodiment, the memory 804 may be a memory component of a reconfigurable logic device, e.g. an FPGA. In one embodiment, the memory 804 includes a cache or random access memory for the processor 802. In alternative embodiments, the memory 804 is separate from the processor 802, such as a cache memory of a processor, the system memory, or other memory. The memory 804 may be an external storage device or database for storing data. Examples include a hard drive, compact disc ("CD"), digital video disc ("DVD"), memory card, memory stick, floppy disc, universal serial bus ("USB") memory device, or any other device operative to store data. The memory 804 is operable to store instructions executable by the processor 802. The functions, acts or tasks illustrated in the figures or

described herein may be performed by the programmed processor **802** executing the instructions **812** stored in the memory **804**. The functions, acts or tasks are independent of the particular type of instructions set, storage media, processor or processing strategy and may be performed by software, hardware, integrated circuits, firm-ware, micro-code and the like, operating alone or in combination. Likewise, processing strategies may include multiprocessing, multitasking, parallel processing and the like.

[0198] As shown, the computer system **800** may further include a display unit **814**, such as a liquid crystal display (LCD), an organic light emitting diode (OLED), a flat panel display, a solid state display, a cathode ray tube (CRT), a projector, a printer or other now known or later developed display device for outputting determined information. The display **814** may act as an interface for the user to see the functioning of the processor **802**, or specifically as an interface with the software stored in the memory **804** or in the drive unit **806**.

[0199] Additionally, the computer system **800** may include an input device **816** configured to allow a user to interact with any of the components of system **800**. The input device **816** may be a number pad, a keyboard, or a cursor control device, such as a mouse, or a joystick, touch screen display, remote control or any other device operative to interact with the system **800**.

[0200] In a particular embodiment, as depicted in FIG. 8, the computer system **800** may also include a disk or optical drive unit **806**. The disk drive unit **806** may include a computer-readable medium **810** in which one or more sets of instructions **812**, e.g. software, can be embedded. Further, the instructions **812** may embody one or more of the methods or logic as described herein. In a particular embodiment, the instructions **812** may reside completely, or at least partially, within the memory **804** and/or within the processor **802** during execution by the computer system **800**. The memory **804** and the processor **802** also may include computer-readable media as discussed above.

[0201] The present disclosure contemplates a computer-readable medium that includes instructions **812** or receives and executes instructions **812** responsive to a propagated signal, so that a device connected to a network **820** can communicate voice, video, audio, images or any other data over the network **820**. Further, the instructions **812** may be transmitted or received over the network **820** via a communication interface **818**. The communication interface **818** may be a part of the processor **802** or may be a separate component. The communication interface **818** may be created in software or may be a physical connection in hardware. The communication interface **818** is configured to connect with a network **820**, external media, the display **814**, or any other components in system **800**, or combinations thereof. The connection with the network **820** may be a physical connection, such as a wired Ethernet connection or may be established wirelessly as discussed below. Likewise, the additional connections with other components of the system **800** may be physical connections or may be established wirelessly.

[0202] The network **820** may include wired networks, wireless networks, or combinations thereof or other device to device interconnection which facilitate the exchange of data there between. The wireless network may be a cellular telephone network, an 802.11, 802.16, 802.20, or WiMax network. Further, the network **820** may be a public network,

such as the Internet, a private network, such as an intranet, or combinations thereof, and may utilize a variety of networking protocols now available or later developed including, but not limited to TCP/IP based networking protocols. The network **820** may further include node-to-node interconnections between computing nodes in a multi-processor computing environment and/or processor interconnections between multiple processing cores of a multi-core processor, such as may be used in either a system-on-chip (SOC) or network-on-chip (NoC) implementation.

[0203] Embodiments of the subject matter and the functional operations described in this specification can be implemented in digital electronic circuitry, or in computer software, firmware, or hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Embodiments of the subject matter described in this specification can be implemented as one or more computer program products, i.e., one or more modules of computer program instructions encoded on a computer readable medium for execution by, or to control the operation of, data processing apparatus. While the computer-readable medium is shown to be a single medium, the term “computer-readable medium” includes a single medium or multiple media, such as a centralized or distributed database, and/or associated caches and servers that store one or more sets of instructions. The term “computer-readable medium” shall also include any medium that is capable of storing, encoding or carrying a set of instructions for execution by a processor or that cause a computer system to perform any one or more of the methods or operations disclosed herein. The computer readable medium can be a machine-readable storage device, a machine-readable storage substrate, a memory device, or a combination of one or more of them. The term “data processing apparatus” encompasses all apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, or multiple processors or computers. The apparatus can include, in addition to hardware, code that creates an execution environment for the computer program in question, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them.

[0204] In a particular non-limiting, exemplary embodiment, the computer-readable medium can include a solid-state memory such as a memory card or other package that houses one or more non-volatile read-only memories. Further, the computer-readable medium can be a random access memory or other volatile re-writable memory. Additionally, the computer-readable medium can include a magneto-optical or optical medium, such as a disk or tapes or other storage device to capture carrier wave signals such as a signal communicated over a transmission medium. A digital file attachment to an e-mail or other self-contained information archive or set of archives may be considered a distribution medium that is a tangible storage medium. Accordingly, the disclosure is considered to include any one or more of a computer-readable medium or a distribution medium and other equivalents and successor media, in which data or instructions may be stored.

[0205] In an alternative embodiment, dedicated hardware implementations, such as application specific integrated circuits, programmable logic arrays and other hardware devices, can be constructed to implement one or more of the

methods described herein. Applications that may include the apparatus and systems of various embodiments can broadly include a variety of electronic and computer systems. One or more embodiments described herein may implement functions using two or more specific interconnected hardware modules or devices with related control and data signals that can be communicated between and through the modules, or as portions of an application-specific integrated circuit. Accordingly, the present system encompasses software, firmware, and hardware implementations.

[0206] In accordance with various embodiments of the present disclosure, the methods described herein may be implemented by software programs executable by a computer system. Further, in an exemplary, non-limited embodiment, implementations can include distributed processing, component/object distributed processing, and parallel processing. Alternatively, virtual computer system processing can be constructed to implement one or more of the methods or functionality as described herein.

[0207] Although the present specification describes components and functions that may be implemented in particular embodiments with reference to particular standards and protocols, the invention is not limited to such standards and protocols. For example, standards for Internet and other packet switched network transmission (e.g., TCP/IP, UDP/IP, HTML, HTTP, HTTPS) represent examples of the state of the art. Such standards are periodically superseded by faster or more efficient equivalents having essentially the same functions. Accordingly, replacement standards and protocols having the same or similar functions as those disclosed herein are considered equivalents thereof.

[0208] A computer program (also known as a program, software, software application, script, or code) can be written in any form of programming language, including compiled or interpreted languages, and it can be deployed in any form, including as a standalone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program does not necessarily correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data (e.g., one or more scripts stored in a markup language document), in a single file dedicated to the program in question, or in multiple coordinated files (e.g., files that store one or more modules, sub programs, or portions of code). A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a communication network.

[0209] The processes and logic flows described in this specification can be performed by one or more programmable processors executing one or more computer programs to perform functions by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated circuit).

[0210] Processors suitable for the execution of a computer program include, by way of example, both general and special purpose microprocessors, and anyone or more processors of any kind of digital computer. Generally, a processor will receive instructions and data from a read only memory or a random access memory or both. The essential elements of a computer are a processor for performing

instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto optical disks, or optical disks. However, a computer need not have such devices. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio player, a Global Positioning System (GPS) receiver, to name just a few. Computer readable media suitable for storing computer program instructions and data include all forms of non-volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto optical disks; and CD ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

[0211] To provide for interaction with a user, embodiments of the subject matter described in this specification can be implemented on a device having a display, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input.

[0212] Embodiments of the subject matter described in this specification can be implemented in a computing system that includes a back end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front end component, e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back end, middleware, or front end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network ("LAN") and a wide area network ("WAN"), e.g., the Internet.

[0213] The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

[0214] The illustrations of the embodiments described herein are intended to provide a general understanding of the structure of the various embodiments. The illustrations are not intended to serve as a complete description of all of the elements and features of apparatus and systems that utilize the structures or methods described herein. Many other embodiments may be apparent to those of skill in the art upon reviewing the disclosure. Other embodiments may be utilized and derived from the disclosure, such that structural and logical substitutions and changes may be made without departing from the scope of the disclosure. Additionally, the

illustrations are merely representational and may not be drawn to scale. Certain proportions within the illustrations may be exaggerated, while other proportions may be minimized. Accordingly, the disclosure and the figures are to be regarded as illustrative rather than restrictive.

[0215] While this specification contains many specifics, these should not be construed as limitations on the scope of the invention or of what may be claimed, but rather as descriptions of features specific to particular embodiments of the invention. Certain features that are described in this specification in the context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable sub-combination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a sub-combination or variation of a sub-combination.

[0216] Similarly, while operations are depicted in the drawings and described herein in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multi-tasking and parallel processing may be advantageous. Moreover, the separation of various system components in the embodiments described above should not be understood as requiring such separation in all embodiments, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

[0217] One or more embodiments of the disclosure may be referred to herein, individually and/or collectively, by the term “invention” merely for convenience and without intending to voluntarily limit the scope of this application to any particular invention or inventive concept. Moreover, although specific embodiments have been illustrated and described herein, it should be appreciated that any subsequent arrangement designed to achieve the same or similar purpose may be substituted for the specific embodiments shown. This disclosure is intended to cover any and all subsequent adaptations or variations of various embodiments. Combinations of the above embodiments, and other embodiments not specifically described herein, will be apparent to those of skill in the art upon reviewing the description.

[0218] The Abstract of the Disclosure is provided to comply with 37 C.F.R. § 1.72(b) and is submitted with the understanding that it will not be used to interpret or limit the scope or meaning of the claims. In addition, in the foregoing Detailed Description, various features may be grouped together or described in a single embodiment for the purpose of streamlining the disclosure. This disclosure is not to be interpreted as reflecting an intention that the claimed embodiments require more features than are expressly recited in each claim. Rather, as the following claims reflect, inventive subject matter may be directed to less than all of the features of any of the disclosed embodiments. Thus, the following claims are incorporated into the Detailed Descrip-

tion, with each claim standing on its own as defining separately claimed subject matter.

[0219] It is therefore intended that the foregoing detailed description be regarded as illustrative rather than limiting, and that it be understood that it is the following claims, including all equivalents, that are intended to define the spirit and scope of this invention.

What is claimed is:

1. A system for processing a data message received via a network from a source, the data message comprising a plurality of portions, including a portion generated by the source dependent upon an event and a portion generated by the source independent of the event, the system comprising:

- a first logic component coupled to the network and operative to receive the data message;
- a second logic component coupled with the first logic component and operative to automatically analyze the received data message to identify the event dependent and event independent portions thereof and further determine if the event independent portion was transmitted by the source prior to occurrence of the event by evaluating whether all of the event independent portion was received prior to any of the event dependent portion; and
- a third logic component coupled with the second logic component and operative to, when it is determined that the event independent portion was transmitted by the source prior to occurrence of the event, take an action with respect to the data message.

2. A system for processing a data message received via a network from a source, the data message comprising a portion that is generated by the source dependent upon an event and a portion that is generated by the source independent of the event, the system comprising:

- a message receiver coupled with the network and operative to receive the data message;
- a message analyzer coupled with the message receiver and operative to automatically analyze the data message to determine if the event independent portion was transmitted by the source prior to occurrence of the event by evaluating whether all of the event independent portion was received prior to any of the event dependent portion;
- a message processor coupled with the message analyzer and operative to, when it is determined that the event independent portion was transmitted by the source prior to occurrence of the event, take an action with respect to the data message.

3. The system of claim 2 wherein the message receiver receives the data message from a network switch.

4. The system of claim 2 wherein the message receiver is comprised by a network switch.

5. The system of claim 2 wherein at least one of the message receiver, message analyzer, or message processor are implemented using at least one field programmable gate array.

6. The system of claim 2 wherein each portion of the received data message includes data indicative of a sequence which is used by the message analyzer to determine the order of the received portions of the data message.

7. The system of claim 2 wherein each data message may comprise data indicative of a message type, at least one message characteristic, and a message payload.

8. The system of claim 7 wherein the message type comprises a TCP/IP message.

9. The system of claim 7 wherein the message type comprises an Ethernet frame.

10. The system of claim 7 wherein the message type comprises a FIX message.

11. The system of claim 7 wherein the at least one characteristic is a size of the data message.

12. The system of claim 2 wherein the source corresponds to at least one of a plurality of market participants.

13. The system of claim 2 wherein the action further comprises disconnecting a connection to the network for the source of the data message when the data message is one of a plurality of received data messages received from the source and a network activity level for the source, as determined by the plurality of received data messages received from the source, exceeds a network activity level threshold as determined by the message analyzer.

14. The system of claim 2 wherein the action further comprises transmitting the data message to a network switch subsequent to taking the action.

15. The system of claim 2 wherein the action further comprises transmitting the data message to a financial market subsequent to taking the action.

16. The system of claim 2 wherein the action further comprises transmitting the data message to a buffer subsequent to taking the action.

17. The system of claim 2 wherein the action further comprises applying a delay to the data message.

18. The system of claim 17 wherein the delay comprises delaying forwarding the message, the delay being computed based on a size for the data message, calculating a product of the size and a serialization time for the data message, and the message processor is further operative to add the product to a time of receipt for the data message.

19. The system of claim 17 wherein the delay comprises delaying forwarding the message, the delay being computed based on a variance in serialization time for the data message and the message processor is further operative to add the variance to a time of receipt for the data message.

20. A method for processing a data message received via a network from a source, the data message comprising a portion that is generated by the source dependent upon an event and a portion that is generated by the source independent of the event, the method comprising:

receiving a data message at a message receiver coupled with the network;

analyzing automatically the data message at a message analyzer to determine if the event independent portion was transmitted by the source prior to occurrence of the event by evaluating whether all of the event independent portion was received prior to any of the event dependent portion; and

processing the data message at a message processor and taking an action with respect to the data message when it is determined that the event independent portion was transmitted by the source prior to occurrence of the event.

21. The method of claim 20 wherein the data message is received from a network switch.

22. The method of claim 20 wherein each portion of the received data message includes data indicative of a sequence.

23. The method of claim 20 wherein the source corresponds to at least one of a plurality of market participants.

24. The method of claim 20 wherein each data message may comprise data indicative of a message type, at least one message characteristic, and a message payload.

25. The method of claim 24 wherein the message type is a TCP/IP message.

26. The method of claim 24 wherein the message type is an Ethernet frame.

27. The method of claim 24 wherein the message type comprises a FIX message.

28. The method of claim 24 wherein the at least one characteristic is a size of the data message.

29. The method of claim 20 wherein analyzing the data message comprises calculating a size for the data message, calculating a product of the size and a serialization time for the data message, and processing the data message further comprises adding the product to a time of receipt for the data message.

30. The method of claim 20 wherein analyzing the data message comprises calculating a variance in serialization time for the data message and processing the data message further comprises adding the variance to a time of receipt for the data message.

31. The method of claim 20 wherein processing the data message includes disconnecting a connection to the network for the source of the data message when the data message is one of many received data messages received from the source and a network activity level for the source, as determined by the many received data messages received from the source, when the network activity level for the source exceeds a network activity level threshold as determined by analyzing the message.

32. The method of claim 20 wherein processing the data message comprises applying a delay to the data message.

33. The method of claim 20 wherein processing the message comprises transmitting the data message to a network switch subsequent to taking the action.

34. The method of claim 20 wherein processing the message comprises transmitting data message to a financial market subsequent to taking the action.

35. The method of claim 20 wherein processing the message comprises adding the data message to a buffer subsequent to taking the action.

36. A system for processing a data message received via a network from a source, the data message comprising a portion that is generated by the source dependent upon an event and a portion that is generated by the source independent of the event, the system comprising:

a means for receiving the data message;

a means for automatically analyzing the received data message to determine if the event independent portion was transmitted by the source prior to occurrence of the event by evaluating whether all of the event independent portion was received prior to any of the event dependent portion; and

a means for, when it is determined that the event independent portion was transmitted by the source prior to occurrence of the event, taking an action with respect to the data message.

* * * * *